NPS ARCHIVE
2000.03
TSAI, P.

DUDLEY KNOX LIBRARY
NAVAL POSTGRADUATE SCHOOL
MONTEREY CA 93943-5101

NAVAL POSTGRADUATE SCHOOL
MONTEREY, CALIFORNIA

THESIS

RE-TARGETING THE GRAZE PERFORMANCE DEBUGGING TOOL FOR JAVA THREADS AND ANALYZING THE RE-TARGETING TO AUTOMATICALLY PARALLELIZED (FORTRAN) CODE

by

Pedro T.H. Tsai

March 2000

Thesis Advisor: Debra Hensgen
Second Reader: Rudy Darken

Approved for public release; distribution is unlimited

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	March 2000	Master's Thesis

4. TITLE AND SUBTITLE: RE-TARGETING THE GRAZE PERFORMANCE DEBUGGING TOOL FOR JAVA THREADS AND ANALYZING THE RE-TARGETING TO AUTOMATICALLY PARALLELIZED (FORTRAN) CODE	5. FUNDING NUMBERS
6. AUTHOR(S) Tsai, Pedro, T.H.	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION/AVAILABILITY STATEMENT:	12b. DISTRIBUTION CODE:
Approved for public release; distribution is unlimited.	

13. ABSTRACT *(maximum 200 words)*

This research focuses on the design of a language-independent concept, Glimpse, for performance debugging of multi-threaded programs. This research extends previous work on Graze, a tool designed and implemented for performance debugging of C++ programs. Not only is Glimpse easily portable among different programming languages, (i) it is useful in many different paradigms ranging from few long-lived threads to many short-lived threads; and (ii) it generalizes the concept of intervals over Graze's original definition. Glimpse's portability has been validated by demonstrating its usefulness in performance debugging of both Java programs as well as automatically parallelized FORTRAN programs.

14. SUBJECT TERMS Performance Debugging, Java Threads, Automatically Parallelized FORTRAN applications			15. NUMBER OF PAGES 134
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18 298-102

Approved for public release; distribution is unlimited

RE-TARGETING THE GRAZE PERFORMANCE DEBUGGING TOOL FOR JAVA THREADS AND ANALYZING THE RE-TARGETING TO AUTOMATICALLY PARALLELIZED (FORTRAN) CODE

Pedro T.H. Tsai
Naval Research Lab, Monterey CA

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

NAVAL POSTGRADUATE SCHOOL
March 2000

DUDLEY KNOX LIBRARY
NAVAL POSTGRADUATE SCHOOL
MONTEREY CA 93943-5101

ABSTRACT

This research focuses on the design of a language-independent concept, Glimpse, for performance debugging of multi-threaded programs. This research extends previous work on Graze, a tool designed and implemented for performance debugging of C++ programs. Not only is Glimpse easily portable among different programming languages, (i) it is useful in many different paradigms ranging from few long-lived threads to many short-lived threads; and (ii) it generalizes the concept of intervals over Graze's original definition. Glimpse's portability has been validated by demonstrating its usefulness in performance debugging of both Java programs as well as automatically parallelized FORTRAN programs.

TABLE OF CONTENTS

I. INTRODUCTION 1

1. Motivation 1
2. Background 4
3. Organization 5

II. RELATED WORKS 7

1. Traditional Approach to Performance Debugging 7
2. Paradyn Parallel Performance Tools 9
3. Pablo 15
4. Delphi 21
5. TimeScan 21
6. Summary 24

III. BACKGROUND ON GRAZE 27

1. Graze Specification Language 27
2. Data Collection Facility 30
3. Visualization Tools 32
4. Summary 35

IV. GLIMPSE: GENERALIZING GRAZE AND APPLYING IT TO MULTI-THREADED JAVA PROGRAMS 37

1. Design Considerations 37
2. Event Collection Code Generation 45
3. Summary 49

V. EXPERIENCES USING GLIMPSE WITH JAVA 51

1. Description of the StopLight Program 51
2. Defining Events and Intervals For the StopLight Program 53
3. Visualizing the Result 57
4. Summary 71

VI. APPLYING GLIMPSE TO AUTOMATICALLY PARALLELIZED FORTRAN PROGRAM 73
1. Parallelizing Programs using OpenMP 73
2. Adapting Glimpse to Monitor Multiprocessing FORTRAN Code 76
3. Results from Monitoring Parallel FORTRAN programs . 79
4. Summary 104
VII. SUMMARY 105
1. Future Work 105
2. Experiences using *Glimpse* with concurrent Java and FORTRAN programs 107
APPENDIX A. OVERVIEW OF JAVA AND JAVA THREAD . . . 109
a. Overview of Java 109
b. Java Threads and Synchronization 110
c. Summary of Java Threads and Synchronization 122
APPENDIX B. GRAMMAR OF THE GLIMPSE SPECIFICATION LANGUAGE 123
a. Grammar of the Glimpse Specification Language 123
b. Keywords for the Glimpse Specification Language 125
APPENDIX C. A NOTE ON EXPERIENCE USING THE FINALIZE FACILITY OF JAVA'S GARBAGE COLLECTION 127
LIST OF REFERENCES 131
INITIAL DISTRIBUTION LIST 133

LIST OF FIGURES

1. Data flow diagram for the Pablo's `Autopilot` decision mechanism. . . . 20
2. Gorge can be used to display the interaction between threads in a message passing application. For this example, event data from a total of 12 threads are shown, with event data from thread 0 plotted at the top of the graph and event data from thread 11 plotted at the bottom of the graph. The connecting lines between threads indicate the previously defined `Transit` interval. Time increases along the x-axis. 33
3. Nibble graphing the number of `Sendmsg` and `Recvmsg` events as a function of time in a message passing application. Time increases along the x-axis. 34
4. Nibble showing the number of `Sendmsg` and `Recvmsg` events in thread 0 and 1, and the differences (the lower line graph) between the number of messages received by thread 1 and number of messages sent from thread 0. Time increases along the x-axis. 35
5. Glimpse's Java utility classes for mapping threads to log files. 41
6. Steps for collecting event data: 1) code generation, 2) compiling the monitored program, 3) loading dynamic library during execution. Arrow indicates dependency at the various stages. 48
7. Plot of `StartMotorArrInt` intervals for experiment one. The horizontal bar indicates the interval between when a car is started and when it reaches the intersection. The time scale increases to the right. There are 13 cars in this test case; they are displayed from top to bottom. . . 59
8. Plots of the `StartMotorArrInt` and `StopAndGo` intervals for experiment one. The StopAndGo intervals are bound by + and `diamond` symbols. Car 1, 3, 5, 9, 10, 12 have `StopAndGo` interval. 60

9. Plot of LargeTimer and SmallTimer events with StartMotorArrInt and StopAndGo intervals in experiment one. The symbol + and symbol x depict when the LargeTimer and SmallTimer objects are instantiated in the StopLight program. The first pair of timers is created (shown at the top of the plot) when the Intersection object is initialized, that is before any car objects are created. 62

10. NoWaitSmallTimer interval is represented by the vertical line connection between the symbol x in the car thread and symbol x in the SmallTimer. There are 7 NoWaitSmallTimer intervals, originating from cars 2, 4, 6, 7, 8, 11, and 13. These intervals represent the new SmallTimer objects created by cars that did not have to stop at the intersection. 63

11. Interval StimerLtChanged is shown as the horizontal bar (between symbol x and symbol diamond) in timer threads. These intervals indicate the small timers that change the traffic light as the result of its timeslice value expiring. The interval LtChangedAndGo is shown as a vertical line from the right edge of the StimerLtChanged interval of the timer thread to the right edge of the StopAndGo interval of the car threads. These intervals represent the notification received by the car threads when the traffic light is changed by the timer threads. 64

12. Plot of StartMotorArrInt intervals for experiment two. The horizontal bar indicates the interval between when the car is started and when it reaches the intersection. The time scale increases to the right. There are 13 cars in this test case; they are displayed from top to bottom. . . 66

13. Plot of StartMotorArrInt and StopAndGo intervals for experiment two. The StopAndGo intervals are bound by + and diamond symbols. Only car 1, 5, 12, and 13 have a StopAndGo interval. 67

14. Plot of LargeTimer and SmallTimer events with StartMotorArrInt and StopAndGo intervals in experiment two. The symbol + and symbol x depict when the LargeTimer and SmallTimer objects are instantiated in the StopLight program. The first pair of timers is created (shown at the top of the plot) when the Intersection object is initialized, before any car objects are created. 68

15. NoWaitSmallTimer interval is represented by the vertical line connection between the symbol x in the car thread and symbol x in the SmallTimer. There are 9 NoWaitSmallTimer intervals, originating from car 2, 3, 4, 6, 7, 8, 9, 10 and 11. These intervals represent the SmallTimer objects created by cars that did not have to stop at the intersection. 69

16. Interval LtimerLtChanged and StimerLtChanged are shown as the horizontal bar (between symbol x and symbol diamond). The interval LtChangedAndGo is shown as a vertical line from the right edge of the LtimerLtChanged or StimerLtChanged interval in the timer thread to the right edge of the StopAndGo interval in the car threads. These intervals represent the notification received by the car threads when the traffic light is changed by the timer threads. 70

17. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the SIMPLE schedule type under a low system load condition. Each event_y denotes the completion of a single loop iteration. The Time increases to the right and the number of counts is indicated on the vertical axis. 83

18. Plots of occurrences of event_y for 4 lightweight processes using the SIMPLE schedule type under a low system load condition (from the same data as in Figure 17). Each event_y denotes the completion of a single loop iteration. The Time increases to the right. 84

19. Plots of the counts of event_Y as a function of time. The plots are for 4 lightweight processes using the SIMPLE schedule type under a low system load condition. The horizontal line near count value 45 shows that the count value for that particular process is not increasing with time. This is an indication that the process is blocked during the duration that count value is not increasing. The Time increases to the right and the number of counts is indicated on the vertical axis. 85

20. Plots of the event_y for 4 lightweight processes using the SIMPLE schedule type under a low system load condition (from the same data as in Figure 19). Each event_y denotes the completion of a single loop iteration. The large gap for process 0 (at the top of the graph) corresponds to the horizontal line of Figure 19. The gap indicates that the corresponding process is blocked during that period of time. 86

21. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using DYNAMIC scheduling under a low system load condition. Each event_y denotes the completion of a single loop iteration. The Time increases to the right and the number of counts is indicated on the vertical axis. 88

22. Plots of event_y for 4 lightweight processes using DYNAMIC scheduling under a low system load condition (from the same data as in Figure 21). Each event_y denotes the completion of a single loop iteration. The Time increases to the right. 89

23. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using DYNAMIC scheduling under a medium system load condition. The horizontal line indicates that one of the processes is temporarily blocked during the execution in the Do loop. The Time increases to the right and the number of counts is indicated on the vertical axis. 90

24. Plots of event_Y for 4 lightweight processes using DYNAMIC scheduling under a medium system load condition (from the same data as in Figure 23). Each event_y denotes the completion of single loop iteration. Process 0 (at the top of the graph) has the most number of event_y, i.e., it executes more iterations of the loop than other processes. Process 3 was blocked for a period of time (as shown by the gap between the event sequences). Consequently, it completes the least number of iterations of the loop. 91

25. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the GSS scheduling under a low system load condition. Each event_y denotes the completion of a single loop iteration. The Time increases to the right and the number of counts is indicated on the vertical axis. 93

26. Plots of event_y for 4 lightweight processes using GSS scheduling under a low system load condition (from the same data as in Figure 25). Each event_y denotes the completion of a single loop iteration. The Time increases to the right. 94

27. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the SIMPLE schedule type under a high system load condition. Each event_y denotes the completion of a single loop iteration. The execution time of this case is significantly longer than the SIMPLE schedule, low system load case (see Figure 17). Three of the processes have finished their loop iterations, this is indicated by vertical lines on the left edge of the graph that reach the count value of 100. The remaining process did not complete its work until much later. The Time increases to the right and the number of counts is indicated on the vertical axis. 95

28. Plots of event_y for 4 lightweight processes using SIMPLE scheduling under a high system load condition (from the same data as in Figure 27). Each event_y denotes the completion of a single loop iteration. The Time increases to the right. The clusters of events on the left edge of the graph show that Processes 0, 1, and 3 have finished their portion of loop iterations, where as the process 2 did not complete its portion of loop until much later, as shown by the cluster of events on right edge of the graph. 96

29. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using DYNAMIC scheduling under a high system load condition. The Time increases to the right and the number of counts is indicated on the vertical axis. Although the workload is not evenly distributed among the processes, even under the high load condition, they finish about the same time. 97

30. Plots of event_y for 4 lightweight processes using DYNAMIC scheduling under a high system load condition (from the same data as in Figure 29). Each event_y denotes the completion of a single loop iteration. Although we request 4 processes, only 3 processes participated in the parallel region due to high system load. 98

31. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the GSS schedule type under a high system load condition. The Time increases to the right and the number of counts is indicated on the vertical axis. Although we request 4 processes, only 3 processes actually participated in the parallel region. The workload is not evenly distributed among the processes but they finish about the same time. 100

32. Plots of event_y for 4 lightweight processes using GSS scheduling under a high system load condition (from the same data as in Figure 31). Each event_y denotes the completion of a single loop iteration. Although we request 4 processes, only 3 processes (0, 1, and 3) actually participated in the parallel region due to high system load. 101

33. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the INTERLEAVE schedule type under a high system load condition. Each event_y denotes the completion of a single loop iteration. The execution pattern of this case is similar to SIMPLE schedule under a high system load condition. (see Figure 27). Three of the processes have finished their loop iterations, this is indicated by vertical lines on the left edge of the graph that reach the count value of 100. The remaining process did not complete it work until much later. The Time increases to the right and the number of counts is indicated on the vertical axis. 102

34. Plots of event_y for 4 lightweight processes using INTERLEAVE scheduling under a high system load condition (from the same data as in Figure 33). Each event_y denotes the completion of a single loop iteration. The Time increases to the right. The clusters of events on the left edge of the graph show that Processes 0, 1, and 2 have finished their portion of loop iterations, where as the process 3 did not complete its portion of loop until much later, as shown by the cluster of events on right edge of the graph. 103

ACKNOWLEDGMENTS

I like to thank Debra Hensgen, this thesis would not have been possible without her guidance and teaching. To my parents, for their support and understanding. Finally, to Paul, who was my brother, mentor, and my best friend, with all my love.

I. INTRODUCTION

This thesis focuses on the design and implementation of a language independent performance debugging suite called Glimpse. Glimpse collects profiling data from the execution of multi-threaded programs and provides visualization tools to help programmers analyze the collected data. In particular, it builds upon the approach taken by Graze [Ref. 1], a performance debugging tool that monitors C++ programs that use the Solaris thread library. Glimpse generalizes Graze, making it both language independent and more functional. A language independent tool is needed to permit monitoring of many threaded programs important to the Navy as well as to all of DoD. Examples of such programs are the Master Environmental Library **(MEL)**, which is written in Java; and the US Navy Numerical Weather Prediction **(NWP)** applications, which are written in *FORTRAN*. **MEL** is a new system being developed by Naval Research Laboratory (NRL) for the purpose of disseminating environmental data over the network. It is important in such a system to dynamically identify performance bottlenecks. NWP applications, on the other hand, are legacy *FORTRAN* code that is being ported to high-end workstations (SGI, DEC/Alpha) and automatically parallized. Here bottlenecks may arise simply from executing the code on a new platform and, to make the job more difficult, the people porting the code are not, in many cases, the original programmers.

1. Motivation

The NWP applications, for example, the Navy Operational Global Atmospheric Prediction System **(NOGAPS)** [Ref. 2] and the Coupled Ocean Atmosphere Mesoscale Prediction System **(COAMPS)** [Ref. 3], predict atmospheric parameters such as winds, temperature, pressure, and precipitation. These predictions are based on the solution of sets of thermodynamic and fluid dynamic equations. By integrating these equations forward in time, NOGAPS and COAMPS are able to predict the state of the atmosphere in the near future. Because NOGAPS and COAMPS are compu-

tationally intensive, and the information they produce is perishable, minimizing the execution time of these codes is paramount. Most NWP applications use some form of parallel computing to improve their performance. Rather than reimplementing these applications in a parallel programming language, programmers often only add compiler directives to legacy code. These directives cause the compiler to automatically parallelize the code. Programmers place these compiler directives in what appears to be the comment sections of the source code, usually near the computationally intensive loop constructs. These directives indicate to the compiler that a particular loop may be safely parallelized. The compiler will then perform data flow analysis and attempt to generate the necessary machine code for the parallel execution of the loop.

MEL [Ref. 4], another example of an application that can benefit from performance monitoring, is an online digital library for environmental data and other resources. Environmental resources are often (i) difficult to locate, (ii) frequently duplicative, (iii) independently defined and formatted, and (iv) accessible only through an interface that is unique to each repository. The objective of MEL is to provide the user with a single interface that they can use to discover, query, retrieve, and order environmental data. Conceptually, the MEL system is analogous to the card catalog in a library. The card catalog enables the user to search the entire holdings of the library by specific criteria such as topics, authors and titles. Using the emerging standards for describing geo-spatial data and contents, MEL provides a digital metadata database for environmental resources. Since it is not practical to replicate the existing worldwide system of distributed repositories by creating a single, massive resource site for environmental information, the designers of MEL instead choose a three-tier client-server architecture. The first tier is the user interface to the MEL system. It consists of customers and standard WEB browsers. The second tier is the MEL access site, which consists of both hardware and software. The MEL access site fulfills the dual role of handling incoming queries and orders from the customers,

as well as matching the queries and orders to MEL software resident on the third tier. The third tier is composed of actual resource site databases, including their extraction and delivery processes. MEL software is installed at each resource site to facilitate interactions between the MEL access site and the resource sites. The MEL Service Architecture (**MSA**) is the software library (API's) developed to implement the multi-threaded servers running on each MEL access site and MEL clients running inside the WEB browser. MSA uses the Common Object Request Broker Architecture (CORBA) framework and communicates over the network via the Internet Inter ORB Protocol (IIOP). Because there are typically many users and only a few MEL access sites, the access site server that handles the user query and order is potentially a bottleneck. A performance bottleneck could also occur between the MEL access site and the resource site servers that process the requests.

MEL is an object-oriented client-server application, where as NOGAPS and COAMPS are scientific *FORTRAN* codes requiring substantial floating-point computation. Although the programming language used to implement MEL is quite different from NOGAPS and COAMPS, they share a common feature in that they both use a threaded programming paradigm. Unlike a sequential program, performance bugs in a multi-threaded program can be difficult to find using the conventional performance analysis tools. Particularly difficult to determine are performance bugs due to thread synchronization and communication. For example, in the automatically parallelized *FORTRAN* code, a computational loop might be distributed across several threads, with each thread working on a portion of the loop. Typically, there is an implicit '*barrier synchronization*' after the end of the parallel loop; a thread that completes its portion of work will wait until all other threads have finished their portion of work. If there is a load inbalance between threads, it is possible that some threads will idle for a long time while waiting for threads that are still working. Another type of performance bug is due to the communication overhead between threads. An idle thread might periodically query other threads to see whether they have completed.

```
final public class GzEvent
{
  static ThreadPool pool;
  final static boolean DISABLE_GLIMPSE=false;
  final static boolean NATIVE_TIMER=false;
  static SystemTimer st;

  static public void init()
  {
    if (DISABLE_GLIMPSE) return;
    pool=new ThreadPool();
    if (NATIVE_TIMER==true) st=new SystemTimer();
  }
  static public void init(String data_dir)
  {
    if (DISABLE_GLIMPSE) return;
    pool=new ThreadPool(data_dir);
    if (NATIVE_TIMER==true) st=new SystemTimer();
  }
  static public void close()
  {
    if (DISABLE_GLIMPSE) return;
    pool.closeAll();
  }
  static public long gettime()
  {
    long ts;
    if (NATIVE_TIMER)
        ts=st.gethrtime_ms();
    else
        ts=System.currentTimeMillis()*1000;
    return ts;
  }
  static public void Recvmsg(int src, int seq)
  {
    if (DISABLE_GLIMPSE) return;
    long ts=GzEvent.gettime();
    ThreadData handle= (ThreadData) pool.getThreadSpecific();
    handle.putout(0);
    handle.puttime(ts);
    handle.putout(src);
    handle.putout(seq);
```

```
    return;
  }
  static public void Sendmsg(int src, int seq)
  {
    if (DISABLE_GLIMPSE) return;
    long ts=GzEvent.gettime();
    ThreadData handle= (ThreadData) pool.getThreadSpecific();
    handle.putout(1);
    handle.puttime(ts);
    handle.putout(src);
    handle.putout(seq);
    return;
  }
  static public void closeThreadData()
  {
    if (DISABLE_GLIMPSE) return;
    pool.closeThreadData(Thread.currentThread());
  }
}
```

The name of the generated event logging class defaults to GzEvent, although the user can specify a different name by passing an optional argument to the **builder_Java** program. Calls to the `SystemTimer` methods to get timestamp values, and to write out the values of variables `src` and `seq` are automatically inserted into the body of the event logging methods, `GzEvent.Recvmsg(src,seq)` and `GzEvent.Sendmsg(src,seq)`. A flag, `NATIVE_TIMER`, can be set to allow the logging methods to use either the standard Java timing method, or the native timing function, if one is provided via the JNI mechanism as discussed in the previous section. To instrument the application, the programmer edits the source code and annotates it using `GzEvent.init()`, `GzEvent.Recvmsg()`, and `GzEvent.Sendmsg()`. [2] Two versions of the `GzEvent.init()` methods are provided–one takes an optional argument to specify the directory to which the log files are to be written. If not specified, the current working directory of the monitored program is used. As stated earlier, the

[2]In the current implementation the programmer must also manually insert a call to the `GzEvent.close()` method. See appendix C on using Java Finalize facility.

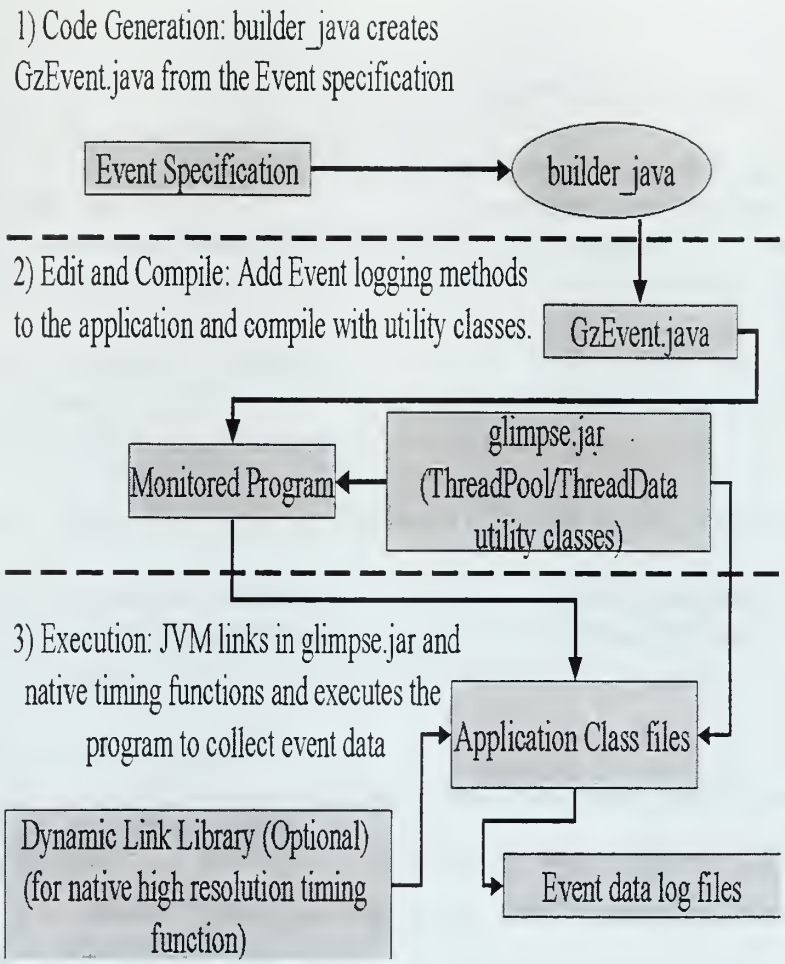

Figure 6. Steps for collecting event data: 1) code generation, 2) compiling the monitored program, 3) loading dynamic library during execution. Arrow indicates dependency at the various stages.

`GzEvent.closeThreadData()` method notifies the ThreadPool object that the calling thread is about to exit and that the log file associated with the exiting thead should be closed. This call is only needed to handle tranisent threads; it prevents data from two different threads being written to the same log file. Finally, `GzEvent.close()` terminates event logging and closes all log files. A flag, `DISABLE_GLIMPSE`, allows the user to turn off the data logging if the application no longer requires performance monitoring.

To compile the monitored program, both the `GzEvent.java` file and `glimpse.jar` utility library need to be supplied to the Java compiler (See Figure 6). The `glimpse.jar` library and dynamic shared library for the native high-resolution timer are then loaded in by the JVM during the execution of the monitored program.

3. Summary

In this chapter, we described the porting issues encountered while adapting *Graze* to monitor multi-threaded Java programs, and the approaches we used to resolve these issues. In particular, the issues with mapping application threads to event log filename, obtaining high-resolution timestamp information at run-time, and generalizing the *interval* definition. In the last section, we descibed how *Glimpse* translates user-defined event specification into the data logging code that is used to instrument the application of interest.

THIS PAGE INTENTIONALLY LEFT BLANK

V. EXPERIENCES USING GLIMPSE WITH JAVA

In this chapter we describe the result from testing *Glimpse* with a multi-threaded Java program. By visualizing the event data collected from execution of the test program and intervals derived from the event data, and by relating these events and intervals to the actions performed by a test program, we demonstrate that *Glimpse* is working properly.

The first section of this chapter describes the structure of a Java example program and the expected interactions between the various threads in this application. The next section describes the event and interval specifications that were used to collect the runtime data, and to visualize that data. Section three describes the data gathered from running the test program with two different inputs. By relating the events and intervals to the expected behaviors of the application, we demonstrated the correctness of the test program and *Glimpse's* Data Collection Facility. In the final section, we summarize the results from testing *Glimpse* with a multi-threaded Java program.

1. Description of the StopLight Program

To evaluate the *Glimpse* software, we applied it to a Java program called *"StopLight"* that simulates traffic flow at an intersection. This program was originally written to show how to use Java threads in a concurrent program. It uses car, intersection, and timer objects to model the flows through the intersection. Each car object is assigned a direction of movement (North-South or East-West) and a departure time. The departure time determines when a car will arrive at the intersection. The timer object controls the traffic lights of the intersection object. Each timer object is initialized with a timeslice value. When that timeslice value expires, the timer object toggles the traffic lights so that the flow direction at the intersection alternates between North-South and East-West. Two pair of timers are used to

control the traffic light at the intersection; their function will be described in next paragraph. Cars traveling in different directions can arrive randomly at the intersection, and more than one car can queue up at the intersection while waiting for the light to change. In this program, a car is not allowed to change direction once it starts running; furthermore there is only one intersection through which all cars must travel.

An intersection should maximize the number of cars passing through, and minimize the time that a car is waiting at the intersection. For example, the green light should stay on longer for the direction with a heavy traffic flow, and should turn red quickly if no more cars are passing through. To implement this behavior in the StopLight program, a timer called LargeTimer (that is a timer with a longer expiration value) and another timer called SmallTimer are used concurrently to control the traffic light. Both timers are started when the intersection object is first initialized. Each time a car passes through the intersection without stopping, the currently active SmallerTimer is stopped and a new one is started. Stopping and restarting this SmallTimer has the effect of prolonging the green light for that given direction. Thus, if a group of cars, traveling along the same direction, arrives at the intersection at close intervals, then the traffic light for that direction will tend to stay green. However, if there is a long gap between cars travelling along the same direction arriving at the intersection, then the SmallTimer will most likely expire before the next car can reach the intersection. When the SmallTimer expires it forces the traffic lights to change, thus allowing the cars traveling along the other direction to proceed. This control logic ensures that cars waiting at the intersection will not wait longer than necessary (not greater than the expiration value of the SmallTimer) if the cross traffic is very sparse.

The LargeTimer ensures that both directions are given their fair share of green light. For example, consider the case where there is a caravan of cars traveling along the North-South direction, while only a few cars are traveling along the East-West

direction. If the gaps between the vehicles in the caravan are very short and there were no LargeTimer object, it is likely then that North-South bound cars could keep extending the traffic light in their favor, and the East-West bound cars would not be able to proceed until all the North-South bound cars were through. The LargeTimer prevents this unfair situation from occurring. Regardless of the number of times that a SmallTimer is restarted, the LargeTimer will eventually expire and force the traffic light to switch. This design guarantees that the maximum waiting time of any cars will not exceed some threshold value, regardless of the uneven traffic distribution between the two flow directions.

In the implementation of the StopLight program, each car and timer object executes as a separate Java thread. Timer threads are assigned higher priority than car threads. The higher priority associated with the timer thread allows it to interrupt other threads in order to change the traffic light at the intersection. The critical sections of this program are associated with changing the traffic light variable of the intersection object. To ensure the synchronized access to the traffic light variable, the intersection object acts as the monitor to the critical sections. (See Appendix A for a discussion on the Java monitor object and synchronization.)

2. Defining Events and Intervals For the StopLight Program

To evaluate the *Glimpse* package with the Stoplight program, we define the following events that are relevant to the car, timer and the intersection objects. For the car objects, we defined three types of events: `StartMotor`, `ArriveIntersection` and `LeaveIntersection`. These events correspond to 1) when the car is started, 2) when it reaches the intersection, and 3) when the car has crossed the intersection. In term of actual location in the `StopLight` program code, these events correspond to 1) when the car object enters its `run()` method, 2) when the car object invokes the `enter()` method of the intersection object, and 3) when the car object returns from the `enter()` method of the intersection object.

For the intersection object we defined the `Stop`, `Go`, and `NoWait` events. The `Stop` event occurs when a car reaching the intersection has to wait for the light to change. The `Go` event occurs when a car waiting at the intersection can proceed. The `NoWait` event occurs when the traffic light direction and car direction are the same, in that case the car just proceeds through the intersection without stopping. The following code shows the locations of event collection points in the `Intersection` object:

```
public class Intersection {
    public int traffic_light ;
    public LargeTimer MajorTimer;
    public SmallTimer MinorTimer;
    . . .
    public synchronized void enter(Car cx, int car_direction)
    {
        /* car going at NORTH or SOUTH direction and so is the current
           traffic light for the intersection. */
        if ( traffic_light == NSDirection &&
             (car_direction == NORTH || car_direction == SOUTH) )
        {
            GzEvent.NoWait(traffic_light);
            MinorTimer.stop();
            MinorTimer = new SmallTimer(200,this);
        }
        /* car going at EAST or WEST direction and so is the current
           traffic light for the intersection. */
        else if ( traffic_light == EWDirection &&
                  (car_direction == EAST || car_direction == WEST) )
        {
            GzEvent.NoWait(traffic_light);
            MinorTimer.stop();
            MinorTimer = new SmallTimer(200,this);
        }
        /* car direction and traffic light direction does not match,
           car must wait until the light changed. */
        else {
            try {
                GzEvent.Stop(traffic_light);
                wait();
```

```
                GzEvent.Go(traffic_light);
            } catch (Exception e) {}
        }
    }
}
```

In the `enter()` method of the Intersection object, the first two '`if`' statement blocks handle the cases where the car direction matches the traffic light direction. In these two `if` statement blocks, the old `SmallTimer` is killed and a new `SmallTimer` is started as per discussion in section one, and execution continues without blocking. In the third `if` statement block, the execution of the `enter()` method is blocked on the call to the `wait()` method. This car will remain blocked until another thread, i.e., the timer thread, acts on the `traffic_light` variable and invokes the `notifyall()` method. (See Appendix A on the `wait()` and `notifyall()` mechanism for synchronizing access to shared variables from multiple threads.)

For the timer objects, we define the `LightChanged` event. This event occurs when the timer acts on the traffic light variable of the Intersection object. We also defined the `SmallTimer` and `LargeTimer` events to indicate when a timer is instantiated. The following code shows the locations where we place the `LightChanged` event in the `LargeTimer` object:

```
public class LargeTimer extends Thread {
    private int timeslice;
    private Intersection in;
    . . .
    public void run()
    {
       try
       {
           /* Sleeping until timeslice expired */
           sleep(timeslice);
           /* Enter critical section and switch the traffic light
              variable of the Intersection object.  Notify any other
              threads waiting on the Intersection object that
              traffic light has changed. */
           synchronized (in)
```

```
            {
                in.MinorTimer.stop();
                if (in.traffic_light==NSDirection) {
                    in.traffic_light = EWDirection;
                    GzEvent.LightChanged(EWDirection);
                    in.notifyAll();
                }
                else
                {
                    in.traffic_light = NSDirection ;
                    GzEvent.LightChanged(NSDirection);
                    in.notifyAll();
                }
                /*  Reset all timers prior to exiting. */
                in.MajorTimer = new LargeTimer(800,in);
                in.MinorTimer = new SmallTimer(200,in);
                GzEvent.closeThreadData();
            }
        } catch (InterruptedException e) { }
    }
}
```

Upon the expiration of the timeslice, the timer (i) toggles the traffic light variable of the Intersection object, (ii) notifies any other threads waiting for the traffic light to change, and (iii) resets the Large and Small timers controlling the Intersection object before exiting.

The event specification for the `StopLight` program is shown below. Additional attribute information is recorded for some of the events. For examples, for events associated with the car objects, the direction of car movement is recorded; for events associated with the intersection and timer objects, the direction of the traffic light is saved to the log files.

```
graze msgs
. . .
event Stop(traffic_light_direction) = symbol(plus);
event Go(traffic_light_direction) = symbol(diamond);
event NoWait(traffic_light_direction) = symbol(x);
event SmallTimer(traffic_light_direction) = symbol(x);
```

```
event LargeTimer(traffic_light_direction) = symbol(plus);
event LightChanged(traffic_light_direction) = symbol(diamond) ;
event ArriveIntersection(car_direction) = symbol(x) ;
event LeaveIntersection(car_direction) = symbol(plus) ;
event StartMotor = symbol(box);
. . .
end msgs.
```

From the event definitions, we specify the following intervals to be constructed and visualized:

```
graze msgs
. . .
interval StartMotorArrInt [StartMotor, ArriveIntersection]=rectangle;
interval ArriveLeave [ArriveIntersection,LeaveIntersection]=rectangle;
interval ArriveAndStop [ArriveIntersection, Stop]=rectangle;
interval StopAndGo [Stop, Go]=rectangle;
interval GoAndLeaveInt [ Go, LeaveIntersection ]  = rectangle;
interval STimerLtChanged [ SmallTimer, LightChanged ] =  rectangle;
interval LTimerLtChanged [ LargeTimer, LightChanged ] =  rectangle;
interval LtChangedAndGo [ LightChanged <- Go] Match = line;
interval NoWaitSmallTimer [NoWait,SmallTimer] Match = line;
. . .
end msgs.
```

`StartMotorArrInt` shows the interval between when a car starts and when it reaches the intersection. The `ArriveLeave` interval depicts cars that arrive at the intersection and go through it without stopping, whereas `ArriveAndStop` interval shows cars that arrive at the intersection and wait for the light to changed. The interval `StopAndGo` shows how long a car spends waiting for the light to change. The `STimerLtChanged` and `LTimerLtChanged` depict the intervals between when the timer is started and when it actually acts to change the traffic light. The `LtChangedAndGo` and `NoWaitSmallTimer` intervals associate events occurring on different threads.

3. Visualizing the Result

We used two different traffic patterns to test the working of Glimpse with the StopLight program. The traffic patterns are configured by setting the departure

time and the direction of the cars. In test case one, cars are initialized with random departure time and different directions. In test case two, a stream of cars traveling along the same direction is started at fairly close intervals. This steady stream of traffic flow is interspersed with a few cars traveling along the other direction. From the previous discussion on the StopLight program, we would expect that in test case one, where cars arrive randomly at the intersection while traveling at different direction, it is more likely that we should see the SmallTimer acts to switch the traffic light. In test case two, a steady flow of cars traveling along the same direction arriving at the intersection will tend to prevent the SmallTimer from changing the traffic light, but will favor the LargeTimer.

We ran the `StopLight` program with the two test cases multiple times on a SGI workstation (with the Java SDK 1.1.6 environment). The results vary slightly from run to run, but the general pattern is consistent between different runs. The following plots show the representative results.

a. Experiment One

In Figure 7, the interval `StartMotorArrInt` are plotted using the *gorge* visualization tool. The symbol x at the right edge of the horizontal bar indicates when a car arrives at the intersection. Because of the different departure times assigned to the cars, car objects instantiated later in the main program can reach the intersection earlier. For example, car 4 reaches the intersection before car 3.

In Figure 8, we added the `StopAndGo` interval to the plot. The `StopAndGo` interval depicts a car that must stop at the intersection for the light to change. From the plot, we see that car 1, 3, 5, 9, 10, 12 stopped at intersection before they were allowed to continue through. Next, we add the events associated with timer objects to the plot. The events `LargeTimer` and `SmallTimer` are depicted by the symbol + and the symbol x respectively. The first occurrence of a pair of `LargeTimer` and `SmallTimer` are associated with the instantiation of the `Intersection` object. These two events are shown at top of the plot (Figure 9) before the horizontal bars

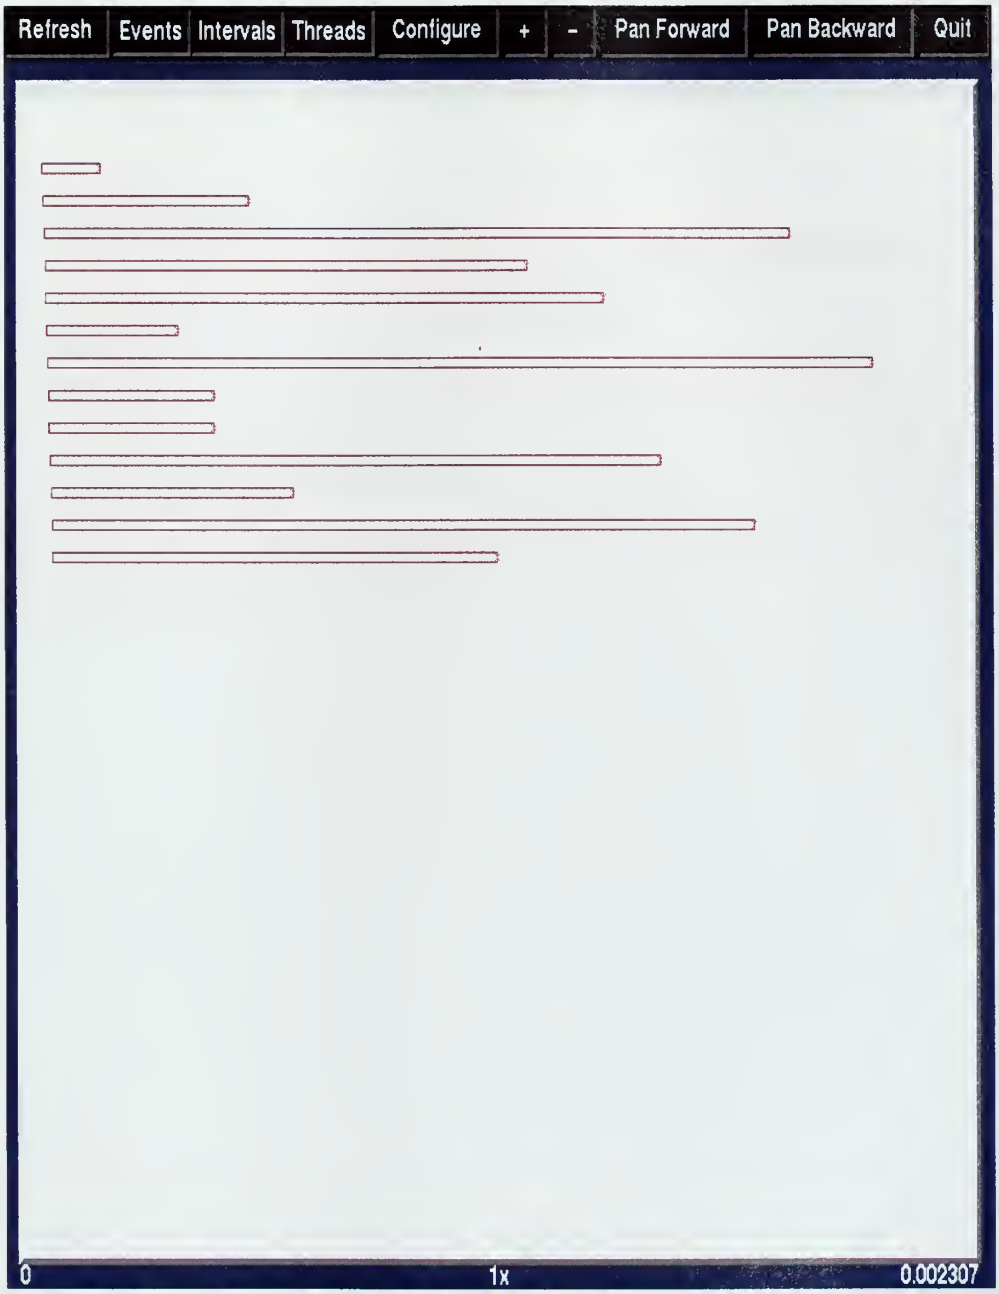

Figure 7. Plot of `StartMotorArrInt` intervals for experiment one. The horizontal bar indicates the interval between when a car is started and when it reaches the intersection. The time scale increases to the right. There are 13 cars in this test case; they are displayed from top to bottom.

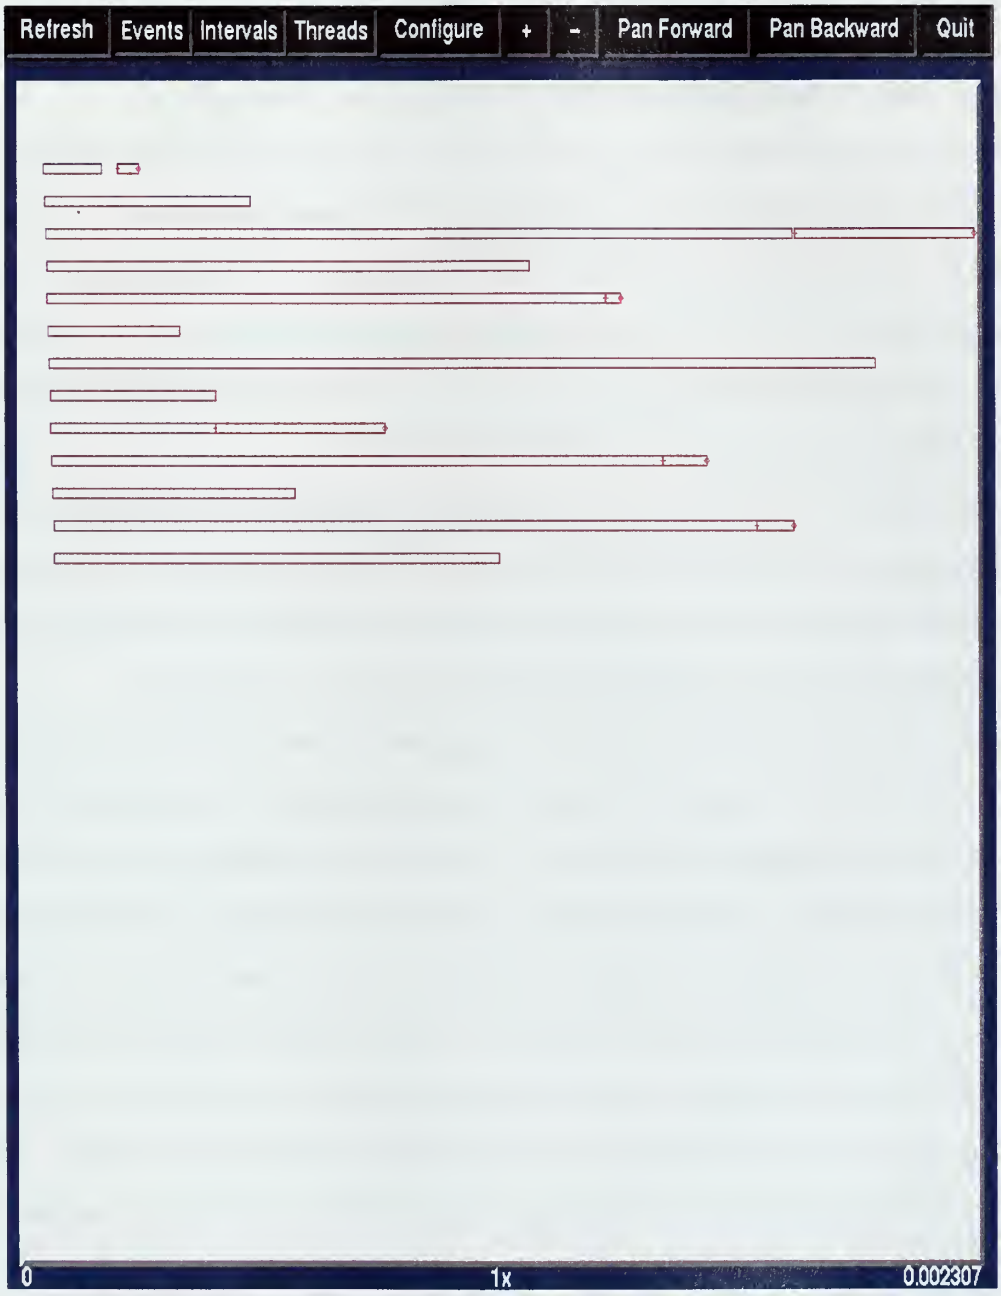

Figure 8. Plots of the `StartMotorArrInt` and `StopAndGo` intervals for experiment one. The StopAndGo intervals are bound by `+` and `diamond` symbols. Car 1, 3, 5, 9, 10, 12 have `StopAndGo` interval.

associated with car objects. Subsequently, we saw more instances of the SmallTimer event further down on the plot. The occurrences of these SmallTimer events are made clear in the next figure.

In Figure 10, we added the interval NoWaitSmallTimer to the plot. The NoWaitSmallTimer interval is represented by the vertical line that connects the right edge of the StartMotorArrInt interval in a car thread to the symbol x denoting the instantiation of a new SmallTimer object in a timer thread. There are 7 NoWaitSmallTimer intervals originating from car 2, 4, 6, 7, 8, 11, and 13; they correspond to cars that did not have to wait at the intersection for the light to change. Recall from the discussion of the *StopLight* program, cars that go through intersection without stopping are extending the duration of the traffic light in their direction by resetting the controlling timers. The interval NoWaitSmallTimer corresponds to this action, and the *gorge* provides a visual representation of this program behavior.

In Figure 11, We add StimerLtChanged and LtChangedAndGo intervals to the plot. The StimerLtChanged represents the interval between when a small timer is created and when its timeslice value has expired, at which time the small timer changes the traffic light direction. The StimerLtChanged interval is depicted as horizontal bars bounded by the x and diamond symbols. These StimerLtChanged intervals are shown below the car threads in Figure 11, except for the first instance of StimerLtChanged interval which is created by the Intersection object. From the plot we see that not every small timer thread contains a StimerLtChanged interval. This is because some small timers are preempted by cars going through the intersection, as discussed in the previous paragraph. The LtChangedAndGo represents the interval between when a car object stops at an intersection and when it starts moving again as the result of the expiring timer changing the traffic light direction. The LtChangedAndGo is depicted as vertical line connecting the right edge of the horizontal bars of the LtChangedAndGo intervals to the right edge of the StopAndGo intervals in the car threads. As can be seen from Figure 11, LtChangedAndGo intervals

Figure 9. Plot of `LargeTimer` and `SmallTimer` events with `StartMotorArrInt` and `StopAndGo` intervals in experiment one. The symbol `+` and symbol `x` depict when the `LargeTimer` and `SmallTimer` objects are instantiated in the `StopLight` program. The first pair of timers is created (shown at the top of the plot) when the `Intersection` object is initialized, that is before any car objects are created.

Figure 10. `NoWaitSmallTimer` interval is represented by the vertical line connection between the symbol `x` in the car thread and symbol `x` in the `SmallTimer`. There are 7 `NoWaitSmallTimer` intervals, originating from cars 2, 4, 6, 7, 8, 11, and 13. These intervals represent the new `SmallTimer` objects created by cars that did not have to stop at the intersection.

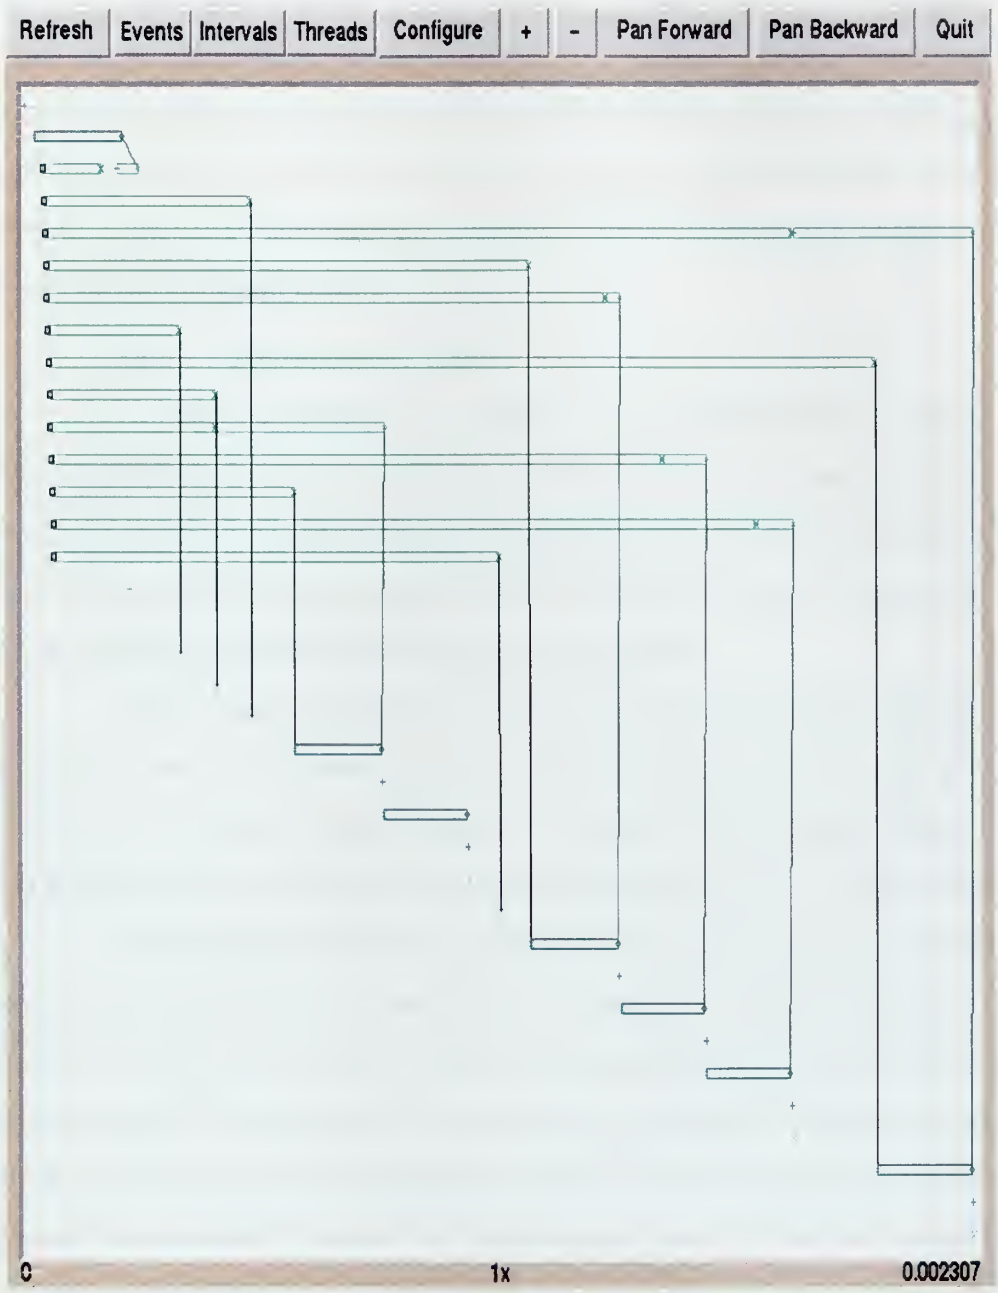

Figure 11. Interval `StimerLtChanged` is shown as the horizontal bar (between symbol `x` and symbol `diamond`) in timer threads. These intervals indicate the small timers that change the traffic light as the result of its timeslice value expiring. The interval `LtChangedAndGo` is shown as a vertical line from the right edge of the `StimerLtChanged` interval of the timer thread to the right edge of the `StopAndGo` interval of the car threads. These intervals represent the notification received by the car threads when the traffic light is changed by the timer threads.

connects car 1, 3, 5, 9, 10, and 12 to `StimerLtChanged` interval. These are cars that were waiting at the intersection for the light to change. The interactions between the timer threads and car threads are consistent with the traffic distribution used in this test case and the design of the *StopLight* program. The visualization of intervals and events allow us to correlate the execution of the application program with the semantics of the program code.

b. Experiment Two

In this experiment, we execute the annotated test program using a different traffic pattern. This traffic pattern consists of cars arriving at the intersection at a constant interval, as shown by the `StartMotorArrInt` intervals in Figure 12. In Figure 13, we added the `StopAndGo` interval to the plot. We see that cars 1, 5, 12, and 13 must stop at the intersection before proceeding.

In the next Figure (Figure 13), the events associated with the creation of `SmallTimer` and `LargeTimer` are shown. These events are associated with the car that did not have to stop at intersection. The interval `NoWaitSmallTimer` in Figure 15 clearly shows the relationship of car threads creating the new timer threads.

The difference between experiment one and experiment two is that in experiment two only one of the `SmallTimer` threads ever reaches expiration. For the experiment two, there is a steady stream of cars going through the intersection, so the exiting `SmallTimer`, except for the last one, is always being re-started. Recall that from the discussion of the *StopLight* program, if cars repeatedly passing through the intersection constantly extend the traffic light, at some point the `LargeTimer` will act to change the traffic light. This is indeed what is shown by the `LtimerLtChanged` interval in Figure 16. The LargeTimer expired and allowed car 5, which has been waiting at the intersection, to proceed. When the `LargeTimer` expired, it created a pair of `SmallTimer` and `LargeTimer`. The event `SmallTimer` and `LargeTimer`, and the interval `StimerLtChanged` are shown at the bottom right of Figure 16. The expiration of the last `SmallTimer` allows cars 12 and 13 to proceed. This is shown by

Figure 12. Plot of `StartMotorArrInt` intervals for experiment two. The horizontal bar indicates the interval between when the car is started and when it reaches the intersection. The time scale increases to the right. There are 13 cars in this test case; they are displayed from top to bottom.

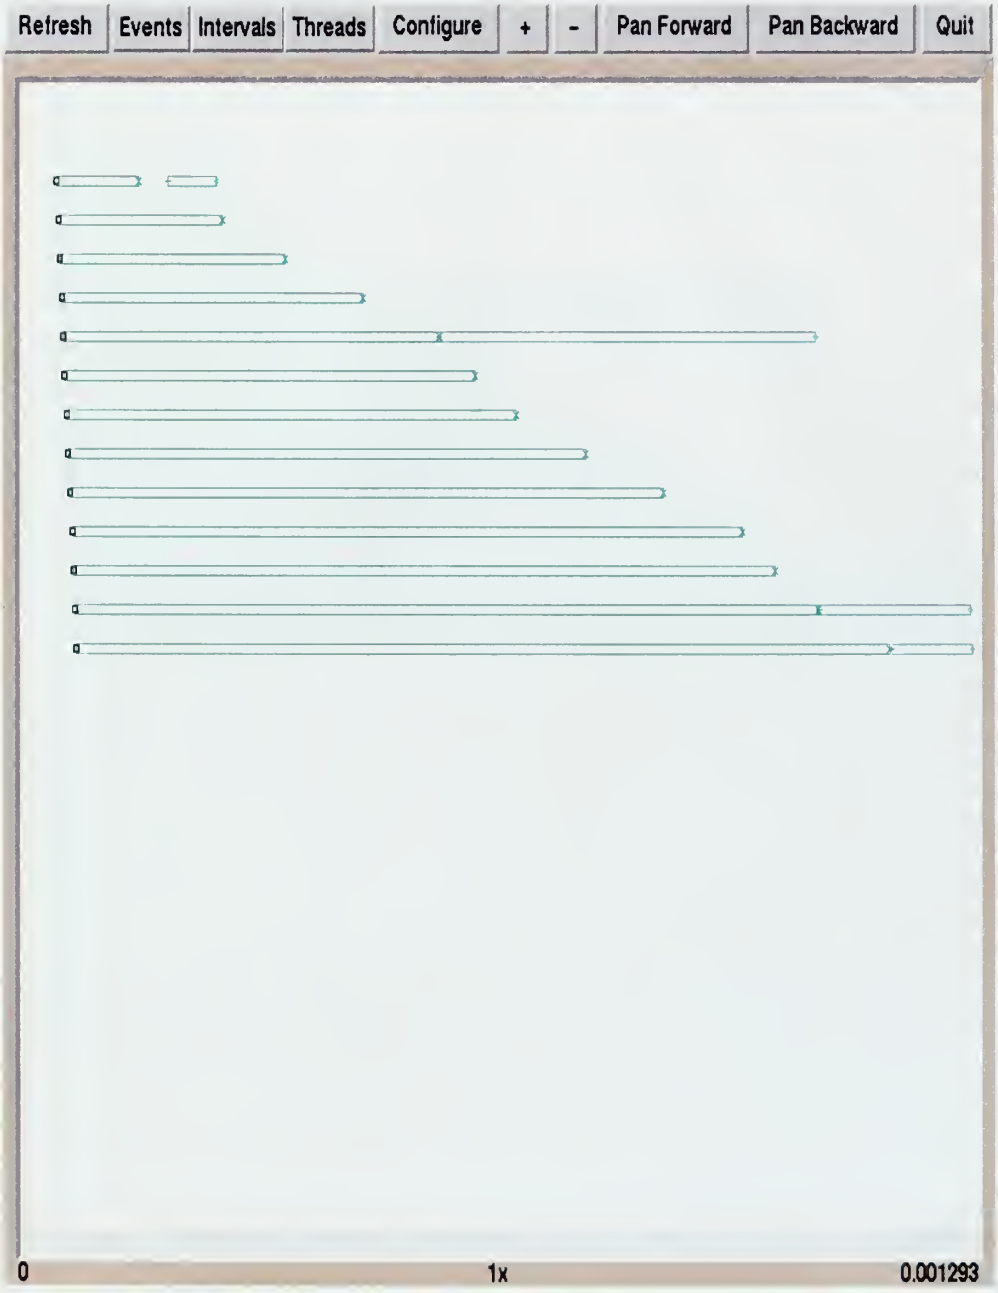

Figure 13. Plot of `StartMotorArrInt` and `StopAndGo` intervals for experiment two. The `StopAndGo` intervals are bound by `+` and `diamond` symbols. Only car 1, 5, 12, and 13 have a `StopAndGo` interval.

Figure 14. Plot of `LargeTimer` and `SmallTimer` events with `StartMotorArrInt` and `StopAndGo` intervals in experiment two. The symbol + and symbol x depict when the `LargeTimer` and `SmallTimer` objects are instantiated in the `StopLight` program. The first pair of timers is created (shown at the top of the plot) when the `Intersection` object is initialized, before any car objects are created.

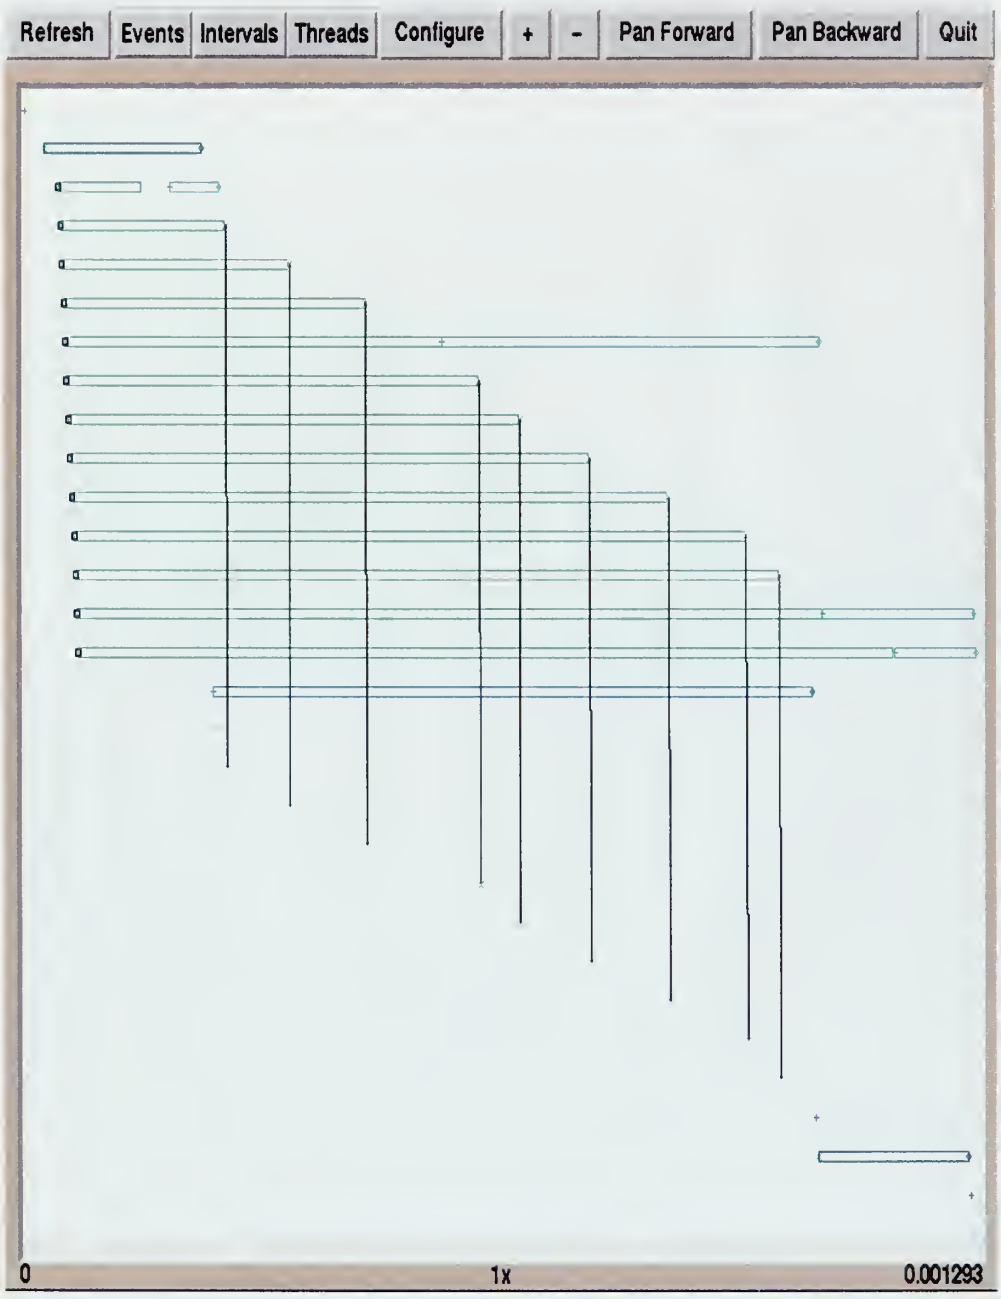

Figure 15. `NoWaitSmallTimer` interval is represented by the vertical line connection between the symbol `x` in the car thread and symbol `x` in the `SmallTimer`. There are 9 `NoWaitSmallTimer` intervals, originating from car 2, 3, 4, 6, 7, 8, 9, 10 and 11. These intervals represent the `SmallTimer` objects created by cars that did not have to stop at the intersection.

Figure 16. Interval `LtimerLtChanged` and `StimerLtChanged` are shown as the horizontal bar (between symbol `x` and symbol `diamond`). The interval `LtChangedAndGo` is shown as a vertical line from the right edge of the `LtimerLtChanged` or `StimerLtChanged` interval in the timer thread to the right edge of the `StopAndGo` interval in the car threads. These intervals represent the notification received by the car threads when the traffic light is changed by the timer threads.

the interval `LtimerLtChanged` that connects the car threads and the last `SmallTimer` thread.

4. Summary

In this chapter, we described a multi-threaded Java program that was used to demonstrate the working of the *Glimpse* Data Collection Facility. We then defined the event and interval specifications that are used to describe the behavior of this test program, in particular the interaction between threads in this application. By running the test program with two different inputs, we collected two experimental data sets. Using the visualization tool *gorge*, we were able to verify that event data collected from the experiments are consistent with the behavior of the application.

THIS PAGE INTENTIONALLY LEFT BLANK

VI. APPLYING GLIMPSE TO AUTOMATICALLY PARALLELIZED FORTRAN PROGRAM

In this chapter, we describe our experiences porting *Glimpse* to monitor the performance of automatically parallelized FORTRAN programs. The FORTRAN programs that we expect *Glimpse* to be useful with are those that use lightweight processes to implement concurrent execution on systems with multiple processors. Although the FORTRAN parallel programming model and the language constructs that FORTRAN uses are quite different from the models and constructs in both Java and C++, we found that the basic approach used *Glimpse* is applicable to the performance analysis of parallel FORTRAN codes. In particular, the language-independent concepts of *events* and *intervals* proved quite useful.

In section one, the OpenMP tool for parallizing FORTRAN code is described. In section two, the modifications we made to *Glimpse's* Data Collection Facility to work with automatically parallelized FORTRAN programs are discussed. In section three, results from testing *Glimpse* with parallel FORTRAN programs are analyzed. In particular we see that different work-load scheduling algorithms have a large impact on the execution time of parallel programs and how *Glimpse* can be useful in understanding this impact. The test programs we used include a simple, loop-level parallel code and an operational weather forecast application. We conclude in section four by describing the experiences that we gained from modifying *Glimpse* to be used with automatically parallelized FORTRAN code and actually using it.

1. Parallelizing Programs using OpenMP

OpenMP [Ref. 18] is a tool from SGI that can be used by FORTRAN and C programmers to allow their code to take advantage of multiple processor environments. In many parallel programming approaches, such as those that use POSIX threads or Message Passing Interface (MPI), programmers must explicitly structure

the program by adding necessary function calls to the source code that handle synchronization, data movement, and other communication between parallel execution tasks. The OpenMP approach, on the other hand, relies on the compiler to handle such details. To set up a program for concurrent execution, the programmer inserts directives for the sections of the code that can safely be executed in parallel. A primary target for these directives is the loop constructs. The OpenMP-enabled compilers process the directives and generate calls to library functions. These library functions create lightweight processes at runtime and distribute loop iterations among them. [1]

The following example (in FORTRAN) illustrates the use of multiprocess directives:

```
      program mploop

      integer maxlen
      parameter(maxlen = 400)
      real*8 a(maxlen), b(maxlen), c(maxlen)
      integer len
      len=maxlen
C
C$DOACROSS SHARE(a,b,c,len),LOCAL(i),MP_SCHEDTYPE=SIMPLE
C$& CHUNK=25
C
      do i = 1, len
          a(i) = b(i) + c(i)
      end do
      stop
      end
```

[1] In this thesis we use the term lightweight process as it is used by SGI. SGI's definition is slightly different from these definitions typically found in textbooks and papers. In their definition, a lightweight process shares its address space with its parent process. In that respect, a lightweight process is similar to a thread. However, their lightweight process has 2 characteristics: (1) it carries a full set of the state information of a process, and (2) the dispatch of their lightweight processes is done in the kernel space. On the Silicon Graphics systems, a lightweight process is created by the **sproc()** system call.

The multiprocess directives are placed in the comment section of the code. They begin with a 'C' in column one and a '$' as the second character. The directives may be followed by optional clauses. If the compiler has multiprocessing turned off, these statements have no effect. This feature allows the identical source code to be compiled for a single processor system, e.g., for debugging purposes, without the multiprocess option. The example demonstrates the use of the `DOACROSS` directive. This directive indicates to the compiler to generate code that will execute multiple iterations of a `DO` loop in parallel. The `C$DOACROSS` directive applies only to the next executable statement that follows, which must be a `DO` loop. The `SHARE` and `LOCAL` clauses specify lists of variables used within the ensuing parallel region. The `LOCAL` clause specifies variables that are local to each parallel process. A variable is declared as LOCAL if its value does not depend on any other iteration of the loop and its value is used only within a single iteration (e.g., the loop-iteration variable `i` in this example). The `SHARE` clause specifies variables that are shared across all parallel processes. If a variable is declared as SHARE, all iterations of the loop use the same copy of the variable. A variable is declared as SHARE if it is only read (not written to) within the loop or if it is an array where each iteration of the loop uses a different element of the array [Ref. 19].

The `MP_SCHEDTYPE` and `CHUNK` clauses affect the way the compiler and run-time environment schedule work among the participating parallel processes in a loop. The default scheduling algorithm (e.g., `SIMPLE`) divides the iterations of the DO loop equally among all participating tasks. Other MP_SCHEDTYPE options are DYNAMIC, INTERLEAVE, and GUIDED SELF-SCHEDULING. These clauses do not affect the correctness of the loop, but can impact the performance of the critical loop. The effect of different work scheduling algorithms will be described in section 3 of this chapter.

The advantage of the compiler-directed approach to automatically parallelizing code is that it hides most of the low-level details associated with parallel program-

ming from the programmers. By adding a few simple directives and hints to the parallel regions of the source code, one can potentially speed up the execution of the application. Additionally, it is simple to change from one type of code partitioning to another. As discussed in section 3 of this chapter the type of scheduling desired can be application specific. Of course, performance bottlenecks will still occur for code parallized by this technique; in particular, in the distribution of workload among the participating tasks and in the interaction with the runtime environment. In the following section, we describe the modifications to *Glimpse* to allow it to monitor such potential performance bottlenecks.

2. Adapting Glimpse to Monitor Multiprocessing FORTRAN Code

Following the approach used for adapting *Graze* to the Java environment, the FORTRAN version of *Glimpse* will use the event and interval specification (Appendix B) and will provide the capability to translate event specifications into event logging subroutines that can be called by the FORTRAN application. The design objectives for the Java implementation (Chapter IV) were also considered in the FORTRAN implementation, that is, (i) maintaining the log file structure and naming convention for compatibility with existing visualization tools, (ii) keeping the number of function calls needed to capture event data to a minimum, and (iii) separating the part of data logging code that is dependent on the event specification from that which is independent in order to improve modularity.

Based on these considerations, we implemented a front-end program, **builder_fortran** and a new utility library, **glimpse_io.a**. The **builder_fortran** program parses the event specifications and generates FORTRAN subroutines. The **glimpse_io.a** library contains utility functions, such as those needed for opening log files, writing to log files, and obtaining timestamp information. The library maintains a list of log file records at runtime, one for each process in the parallel FORTRAN application. A log file record is a C language `struct` that contains information such

as the name of the log file and the file pointer to the log file, as well as a pointer to an internal buffer in which to store the event data. During the execution of parallel FORTRAN programs, the event logging subroutine looks up the appropriate log file record assigned to the given execution process and writes the event data to that buffer. When the event buffer is full, the library writes out the data to the appropriate log file. Unlike the Java programs where some threads might be short lived, the participating processes in the parallel FORTRAN code exists for the duration of the program. [2] Consequently, the method for looking up the appropriate log file record for a given process is straightforward. For example, for a FORTRAN program running with 4 processes, the first process is associated with the element 0 of the log file record array and data from that process are written to 'log.0.'

The utility functions in **glimpse_io.a** library are implemented in the C language. To facilitate implementation of FORTRAN data logging subroutines, the FORTRAN interface to the library is provided by the following functions: `GLIMPSE_INITIALIZE`, `GLIMPSE_CLOSE`, `LOG_EVENT`, and `LOG_EVENT_ATTR`. `GLIMPSE_INITIALIZE` function creates the necessary log file records and initializes the log files. The `GLIMPSE_CLOSE` function flushes out any event data in the memory buffer before closing the log files. The `LOG_EVENT` function writes out the event id and its timestamp to the specified buffer and the `LOG_EVENT_ATTR` function writes out the event attribute. Consider the following event.

```
event write_data(size)
```

The **builder_fortran** program generates the following FORTRAN subroutines:

---

[2]For example, on the Silicon Graphics systems, the main process creates a group of lightweight processes at the start of the program. When the main thread of execution reaches a parallel region in the code, i.e., the `Do` loop, these lightweight processes are unblocked and become active participants in the execution of the `Do` loop. After leaving the parallel region, the lightweight processes become blocked again until the main execution thread reaches another parallel region. The number of lightweight process, a user controllable parameter, is set prior to the program execution. It cannot be changed once the program execution starts.

```
c
c  File: fglimpse.f
c
c  Subroutine to intialize the glimpse io package
c
      subroutine glimpse_init()
         integer num_of_proc(1)
         num_of_proc(1)=mp_numthreads()
         call GLIMPSE_INITIALIZE(num_of_proc)
      return
      end
c
c  Subroutine to close the glimpse io package
c
      subroutine glimpse_terminate()
         call GLIMPSE_CLOSE()
      return
      end
c
c  Subroutine to log event
c
      subroutine write_data(size)
         integer tid(1)
         integer event_id(1)
         integer event_attr(1)
         tid(1)=mp_my_threadnum()
         event_id(1)=0
         call LOG_EVENT(tid,event_id)
         event_attr(1)=size
         call LOG_EVENT_ATTR(tid,size)
      return
      end
```

The FORTRAN subroutine, `glimpse_init()`, initializes the Data Collection Facility by calling the C-function `GLIMPSE_INITIALIZE` to create the necessary log file records and initialize the log files. The function `mp_numthreads()` returns the number of processes that will be used to execute the parallel FORTRAN code. This number is passed to GLIMPSE_INITIALIZE during the initialization. The subroutine `glimpse_terminate` calls the C-function GLIMPSE_CLOSE to flush out any event

data in the memory buffer to files and shutdown the Data Collection Facility. The FORTRAN subroutine for logging the `write_data` event calls the utility functions LOG_EVENT and LOG_EVENT_ATTR to record the event id, timestamps, and event attribute information. The function `mp_my_threadnum()` is called inside the event logging subroutine to identify the process in which each particular event occurs.

The `mp_my_threadnum()` and `mp_numthreads()` are functions in the multitasking library provided by the system vendor. Compilers that process the multiprocess directives use functions from this library to implement the parallel execution of FORTRAN programs. By isolating these system specific calls to the FORTRAN component of the Data Collection Facility, only the **builder_fortran** program needs to be modified to generate the appropriate calls for another vendor library if we move to another platform. The utility functions in the **glimpse_io.a** library are platform independent.

3. Results from Monitoring Parallel FORTRAN programs

In this section we describe our experiences monitoring automatically parallelized FORTRAN programs using the modified *Glimpse* software. We concentrate on how the workload (e.g., the loop iterations) is divided among the participating processes in the parallel region of the code. Before discussion the results, we will first describe the algorithms that the runtime multiprocess library uses to divide up the loop iterations. We recall from the discussion in section one that there is a 'MP_SCHEDTYPE' clause in the compiler directive that can be set, by the user, to guide the scheduling in parallel region. The allowable values of 'MP_SCHEDTYPE' and their meanings are list below [Ref. 18]:

Workload Schedule	Meaning
SIMPLE	Each process executes N/P consecutive iterations starting at Q * (N/P) iterations. The first process to finish takes the remaining chunk, if any.
DYNAMIC	Each process executes C iterations of the loop, starting with the next undone chunk unit, returning for another chunk until none are left.
INTERLEAVE	Each process executes C iterations at (C * Q) + base, (C * 2Q) + base, (C * 3Q) + base, etc. The base = 1, 1 + (C * P), 1 + (C * 2P), etc.
GSS	Guided Self-Scheduling, each process executes chunks of decreasing size, first N/2P, then N/4P, etc

The variables used in the Table are as follows:

N-Number of iterations in the loop, determined from the source or at run time.

P-Number of participating lightweight processes, set by the default or by the user via an environment variable.

Q-Unique identification of a lightweight process, in the range from 0 to P-1.

C-"Chunk" size, set by compiler directive or by environment variable.

The SIMPLE method divides the loop iterations among processes by dividing them into contiguous pieces and assigning one piece to each process. For example, if N is 500, and the number of processes P is 4, then process 0 is assigned loop iterations 1 to 125, process 1 is assigned loop iterations 126 to 250, etc.

In the Dynamic method, the iterations are divided into pieces. The size of a piece is specified with the CHUNK clause. As each process finishes a piece, it grabs the next available piece. For example, if N is 500, P is 4, and CHUNK size is 25, then a possible scenario for process 0 can have it works on iteration 1-25, 101-125, 151-175, 201-225, 301-325, 401-425, 476-500, etc. When a process runs on a CPU

without much competition for CPU cycles, it may process more than an equal share of chunks.

The `Interleave` method also breaks the iterations into pieces of size specified by the CHUNK variable. Execution of those pieces is interleaved among the processes. For example, for P=4, N=500, C=25, process 0 would work on iterations 1-25, 101-125, 201-225, 301-325, 401-425; and process 1 works on iteration 26-50, 126-150, 226-250, etc. Although the number of iterations executed by each process is the same for SIMPLE and INTERLEAVE, INTERLEAVE can outperform SIMPLE when the amount of work to be done in an iteration is a function of iteration number.

In the `GSS` method, the size of pieces varies depending on the number of iterations remaining. In the beginning, larger pieces are parceled out, and as the work nears completion, smaller pieces are parceled out.

We use the example program (see below) from section one to demonstrate the value of the modified *Glimpse* with automatically parallelized FORTRAN code. Two data logging calls, `event_x()` and `event_y()`, denote the entry and exit events for each iterations inside the parallel region. A call to glimpse_init() initializes the Data Collection Facility, and a call to glimpse_terminate() flushes event data in the memory buffers to the log files.

```
      program mploop

      integer maxlen
      parameter(maxlen = 400)
      real*8 a(maxlen), b(maxlen), c(maxlen)
      integer len
      len=maxlen
C
C  Initialize glimpse Data Collection Facility
C
       call glimpse_init()
C
C$DOACROSS SHARE(a,b,c,len),LOCAL(i),MP_SCHEDTYPE=SIMPLE
C$& CHUNK=25
C   MP_SCHEDTYPE=DYNAMIC
```

```
C   MP_SCHEDTYPE=INTERLEAVE
C   MP_SCHEDTYPE=GSS
C
      do i = 1, len
          call event_x()
          a(i) = b(i) + c(i)
          call event_y()
      end do
C
C   Close glimpse Data Collection Facility
C
      call glimpse_terminate()
      stop
      end
```

To evaluate the effect of various workload schedules, we ran the example program multiple times with each schedule mode under different load condition. The number of processes, P, is set to 4, the number of iterations, N, is 400, and the chunk size, C, is 25, for all runs. The example program was run on a 6-processor Silicon Graphics R12000 platform, under various load conditions. The load of the system during the execution was recorded using the UNIX 'w' command. The load on the system gives an indication of the average number of jobs in the CPU queues. For the purpose of this discussion, we consider a load value of less than 1 as low, a value between 1 and 3 as medium, and a value of 4 or greater as high.

Figures 17 and 20 show representative results for the example program running with the SIMPLE workload schedule under a low load condition. Each participating process grabs an equal share of workload and they all take about the same amount of time to complete their portion of the workload. This is shown by all 4 processes reaching the event count of 100 at about the same time (Figure 17), and none of the participating processes experience any delay (Figure 18). In some instances, one of the processes is temporarily blocked. This is shown by the horizontal line in Figure 19 and the corresponding gap in the event sequences in Figure 20. The temporarily blocked processes took longer to finish. Because of the synchronization barrier at the

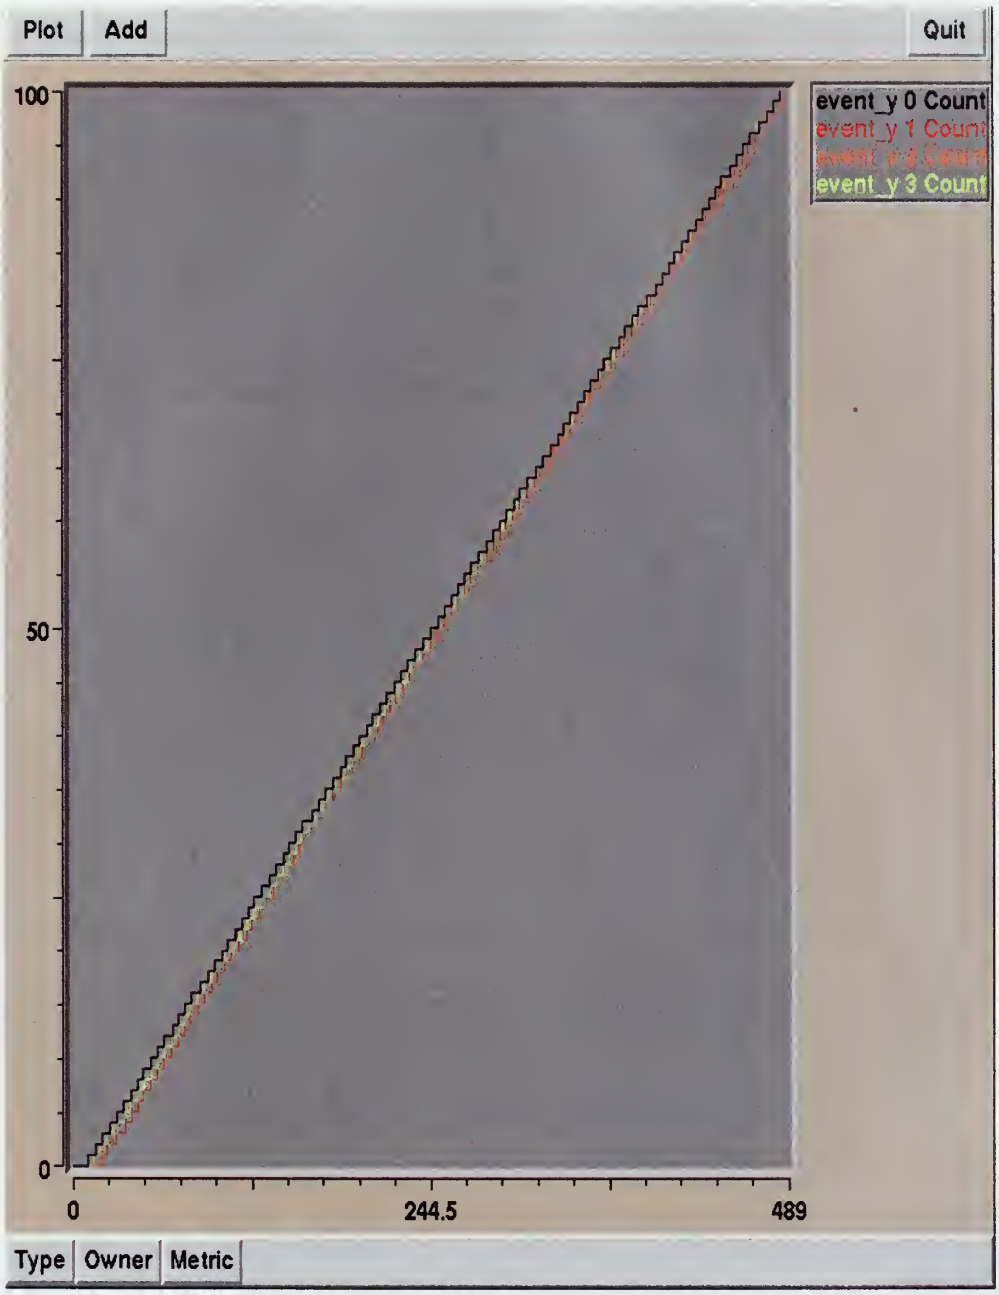

Figure 17. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the SIMPLE schedule type under a low system load condition. Each event_y denotes the completion of a single loop iteration. The Time increases to the right and the number of counts is indicated on the vertical axis.

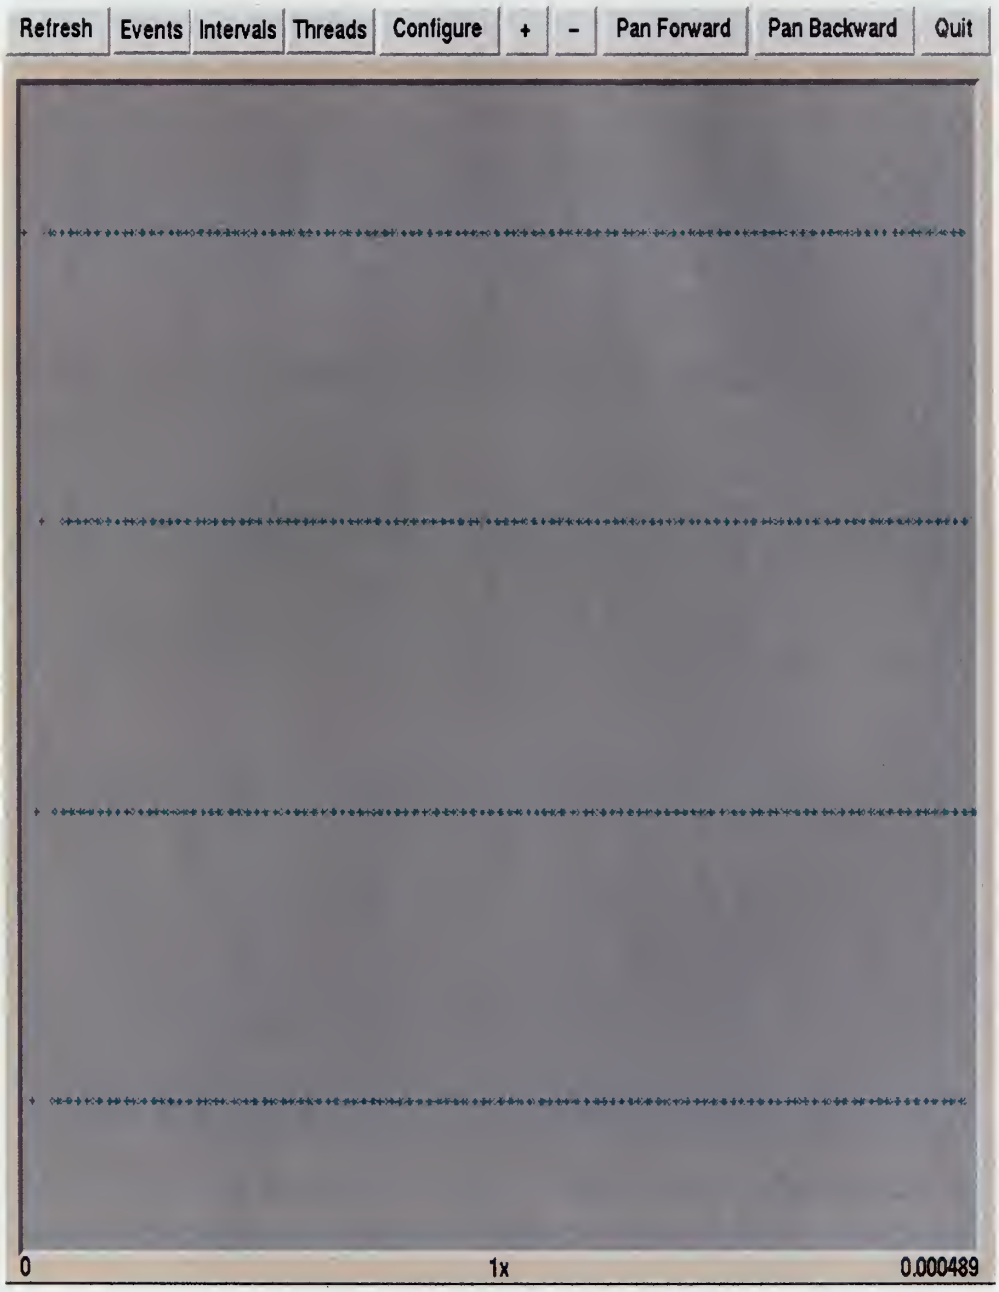

Figure 18. Plots of occurrences of event_y for 4 lightweight processes using the SIMPLE schedule type under a low system load condition (from the same data as in Figure 17). Each event_y denotes the completion of a single loop iteration. The Time increases to the right.

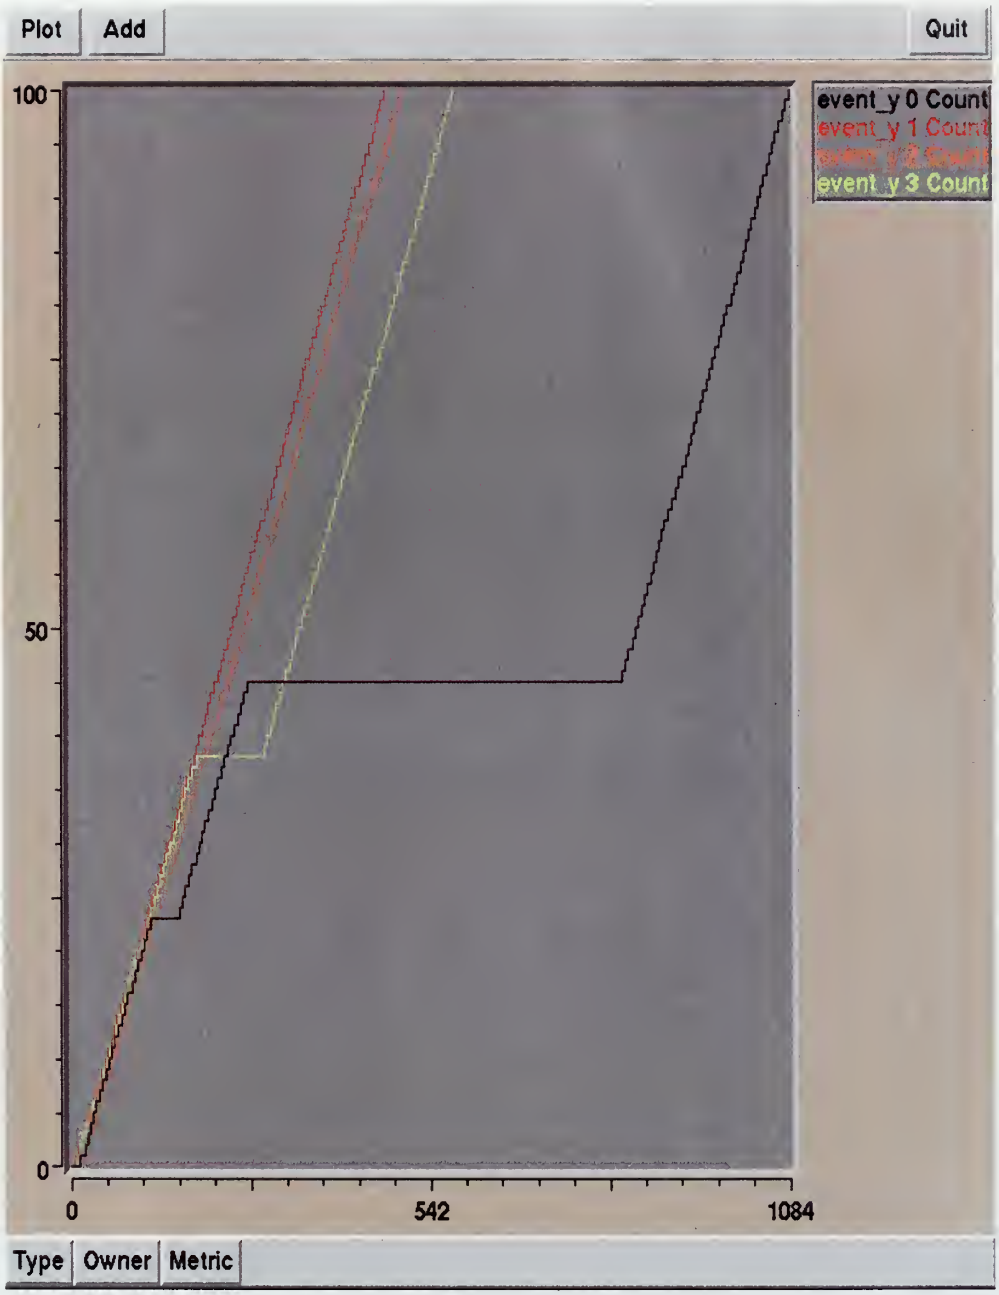

Figure 19. Plots of the counts of event_Y as a function of time. The plots are for 4 lightweight processes using the SIMPLE schedule type under a low system load condition. The horizontal line near count value 45 shows that the count value for that particular process is not increasing with time. This is an indication that the process is blocked during the duration that count value is not increasing. The Time increases to the right and the number of counts is indicated on the vertical axis.

Figure 20. Plots of the event_y for 4 lightweight processes using the SIMPLE schedule type under a low system load condition (from the same data as in Figure 19). Each event_y denotes the completion of a single loop iteration. The large gap for process 0 (at the top of the graph) corresponds to the horizontal line of Figure 19. The gap indicates that the corresponding process is blocked during that period of time.

end of the loop, the completed processes must wait until the last process finishes. As a result, the overall execution time is longer in cases where blocking occurs in at least one process.

We next look at the example program running with the DYNAMIC workload schedule under low to medium load conditions. Figure 21 and 22 show results under a low load condition. The pattern is very similar to the SIMPLE schedule type in that each participating process grabs an equal share of the workload and they all require approximately the same amount of time to complete their portion of the workload. However, runs using the DYNAMIC schedule mode seem to take slightly longer time to complete than runs using the SIMPLE schedule mode under similar load conditions. This is probably due to the higher overheads associated with dynamic scheduling, given that the loop calculation time is very small.

Next results from running under a medium load condition with DYNAMIC scheduling are shown in Figure 23 and 24. Here, the number of iterations is not equally distributed across the processes. Process 3 was temporarily blocked, so it completes fewer iterations than the other processes. All of the processes finished at about the same time. As would be expected, the time it takes to complete the loop using DYNAMIC scheduling type under a low to medium load condition is, on average, slightly longer than the SIMPLE schedule type under a low load condition.

Next, we examine results using the GSS and INTERLEAVE scheduling policies. Under a low to medium load condition, the example program running with a GSS scheduling policy (Figures 25 and 25) exhibits similar behavior to that of the DYNAMIC scheduling policy. The workload across the processes is not equally distributed as in the DYNAMIC case, when a process finishes its chunk of workload, it grabs the next available chunk, so some processes perform more iterations than others. For the INTERLEAVE schedule type, the workload distribution across all of process is the same under our example (we chose the total number of loop iterations and chunk size such that each process gets 4 chunks, each of size 25). The average

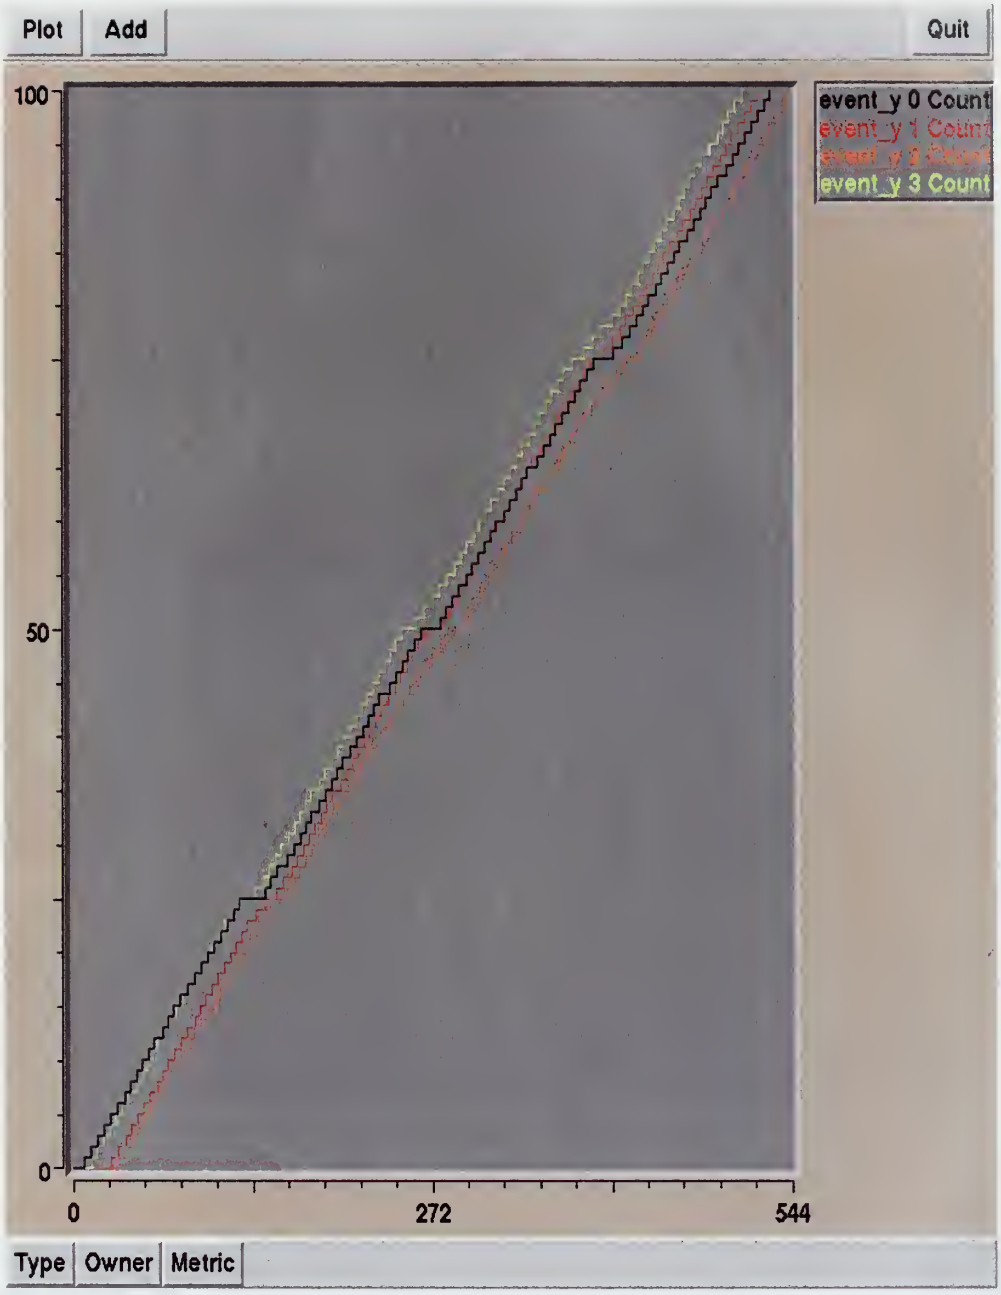

Figure 21. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using DYNAMIC scheduling under a low system load condition. Each event_y denotes the completion of a single loop iteration. The Time increases to the right and the number of counts is indicated on the vertical axis.

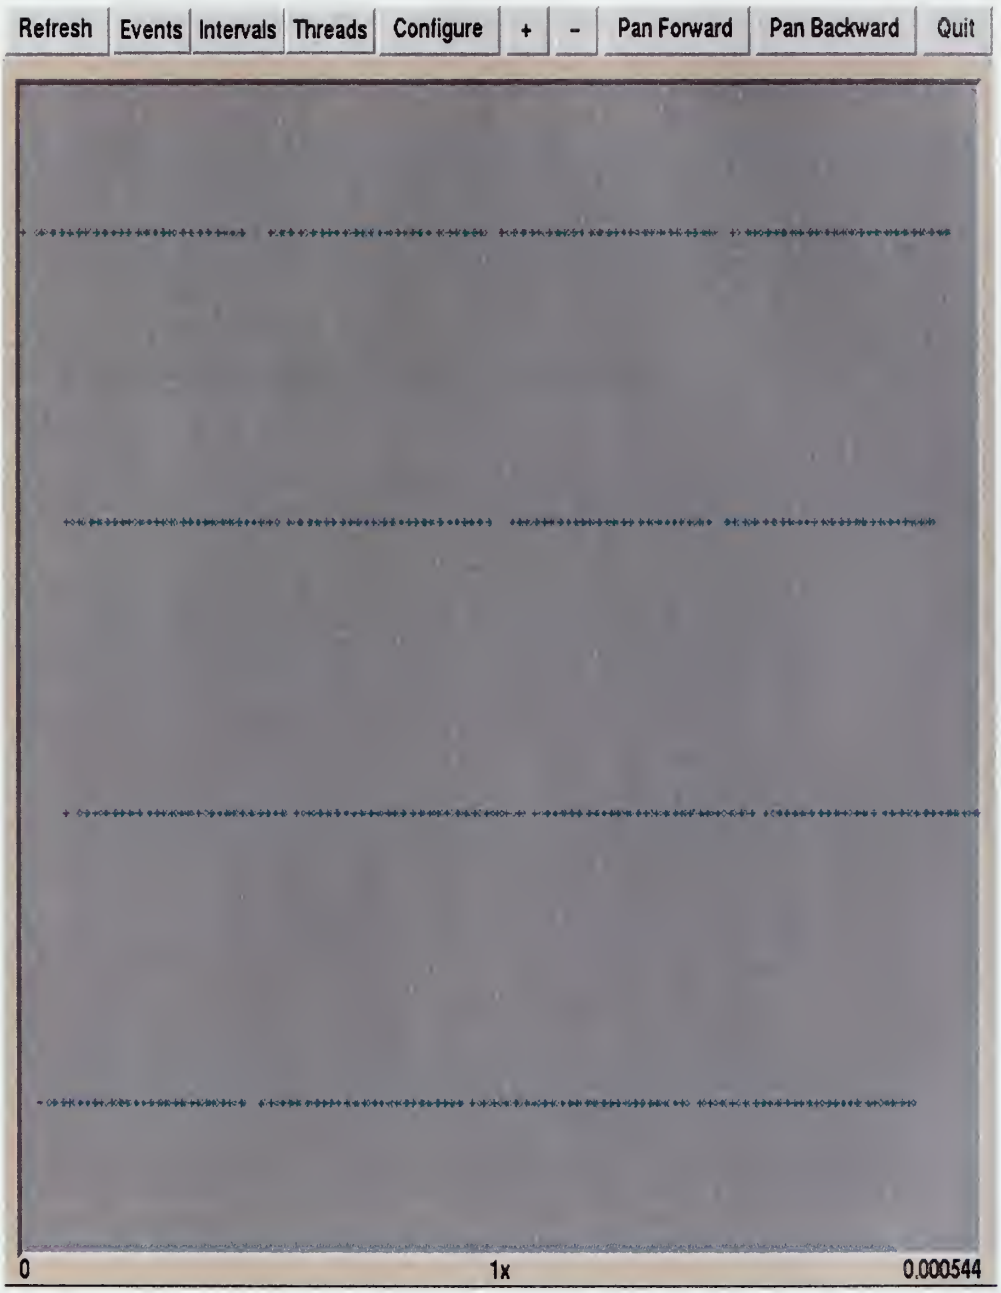

Figure 22. Plots of event_y for 4 lightweight processes using DYNAMIC scheduling under a low system load condition (from the same data as in Figure 21). Each event_y denotes the completion of a single loop iteration. The Time increases to the right.

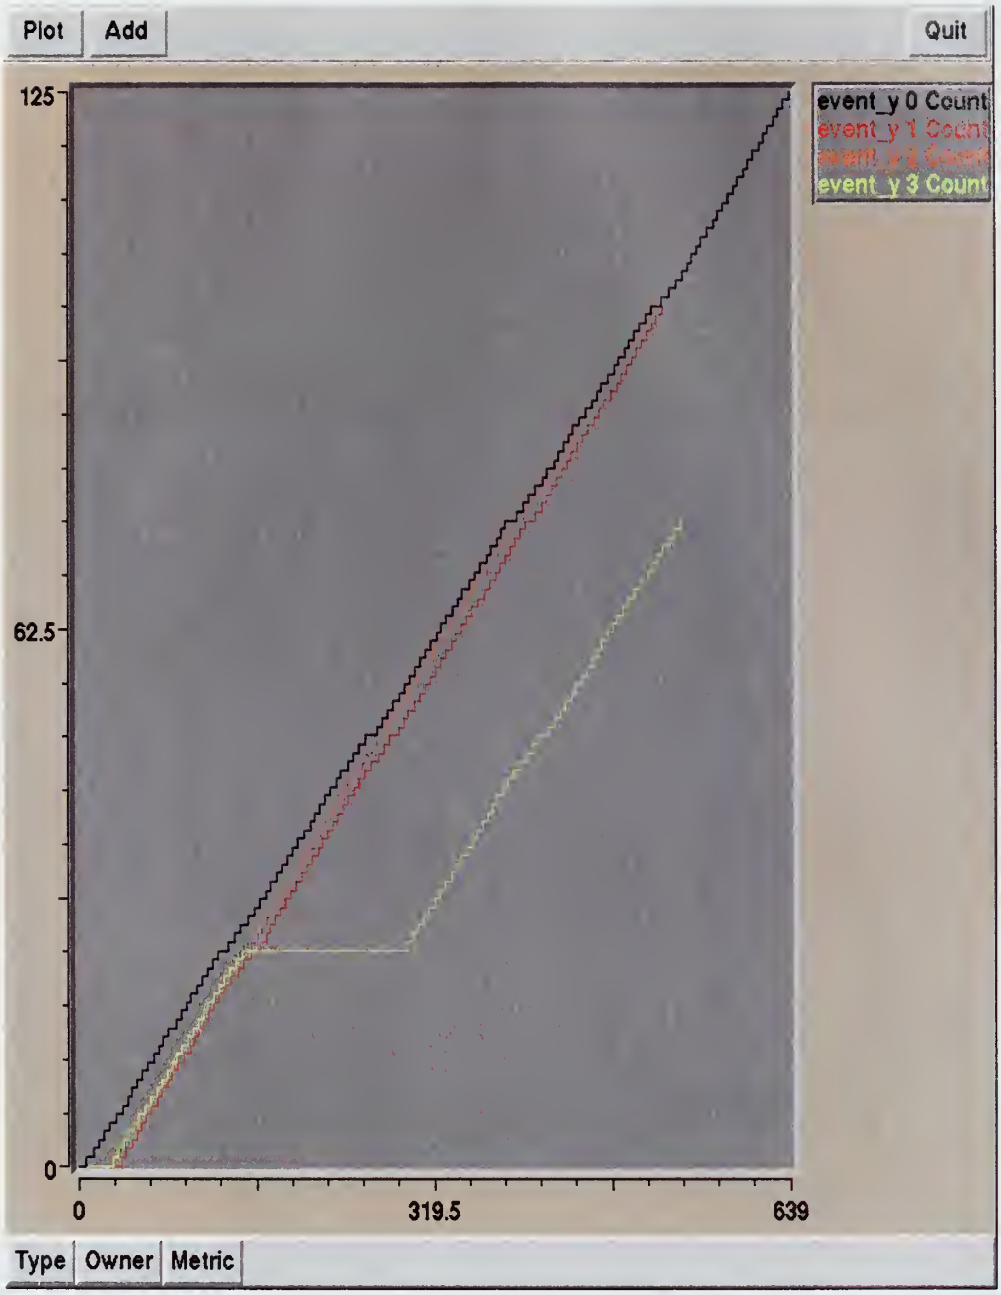

Figure 23. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using DYNAMIC scheduling under a medium system load condition. The horizontal line indicates that one of the processes is temporarily blocked during the execution in the Do loop. The Time increases to the right and the number of counts is indicated on the vertical axis.

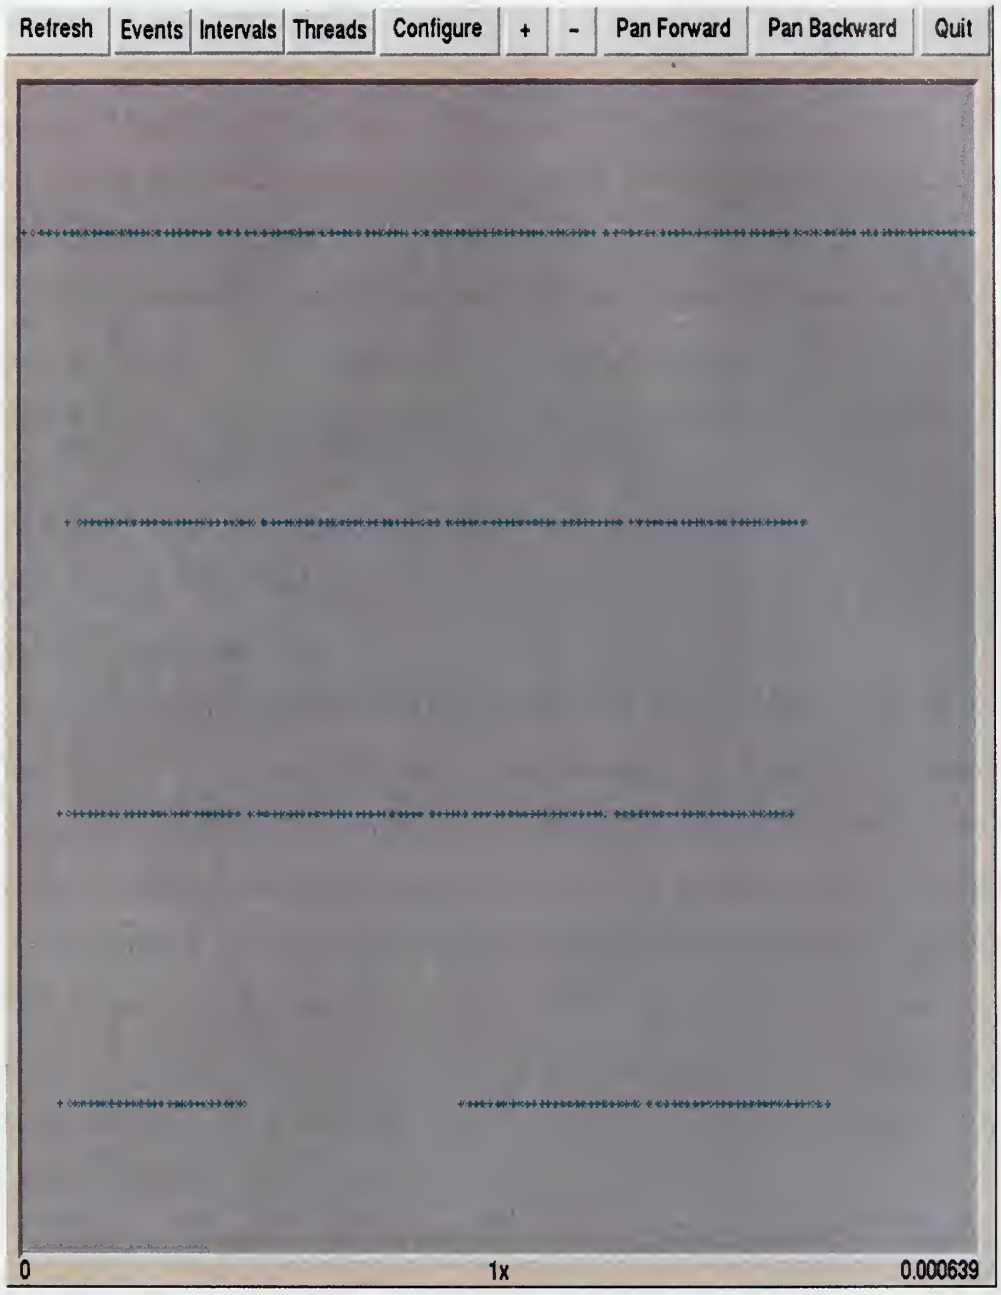

Figure 24. Plots of event_Y for 4 lightweight processes using DYNAMIC scheduling under a medium system load condition (from the same data as in Figure 23). Each event_y denotes the completion of single loop iteration. Process 0 (at the top of the graph) has the most number of event_y, i.e., it executes more iterations of the loop than other processes. Process 3 was blocked for a period of time (as shown by the gap between the event sequences). Consequently, it completes the least number of iterations of the loop.

time to finish the loop is about the same as with the SIMPLE schedule type. This result is expected because the purpose of interleaving iterations across the loop is to improve the access to shared arrays via cache optimization. Since our example uses small arrays (three 400 element arrays of double precision values), the cache optimization is not an issue.

We repeated the experiment with different schedule types by running the same example program under a high load condition. Figures 27 and 28 show the result for the SIMPLE schedule type under a high load condition. The execution time is significantly longer than under the low system load condition. In Figure 28 three of the processes have completed their work, but are forced to wait for the remaining process to finish. The waiting time accounts for most of the execution time for this program to complete.

Next, we look at the DYNAMIC schedule type under a high load condition (Figures 29 and 30). The plots show that although we request 4 processes, only 3 processes participate in the parallel region due to the high system load. The workload is not equally distributed among the 3 processes and all of them finish about the same time. The time to complete the loop for the example program using the DYNAMIC scheduling policy under the high load condition is longer than under the low load condition. However, it does not suffer from the significant performance degradation as experienced by the example program using a SIMPLE scheduling policy under a high load condition.

Next, we examine the results for GSS and INTERLEAVE scheduling types under a high load condition. The example program running with a GSS scheduling policy under high load condition (Figure 31 and 32) exhibits similar characteristics to that of running with the DYNAMIC scheduling policy. Like the previous case, although 4 processes were requested, only 3 processes participated in the parallel region due to high system load. The workload is not equally distributed among the 3 processes, but they all finish about the same time. The performance of the example

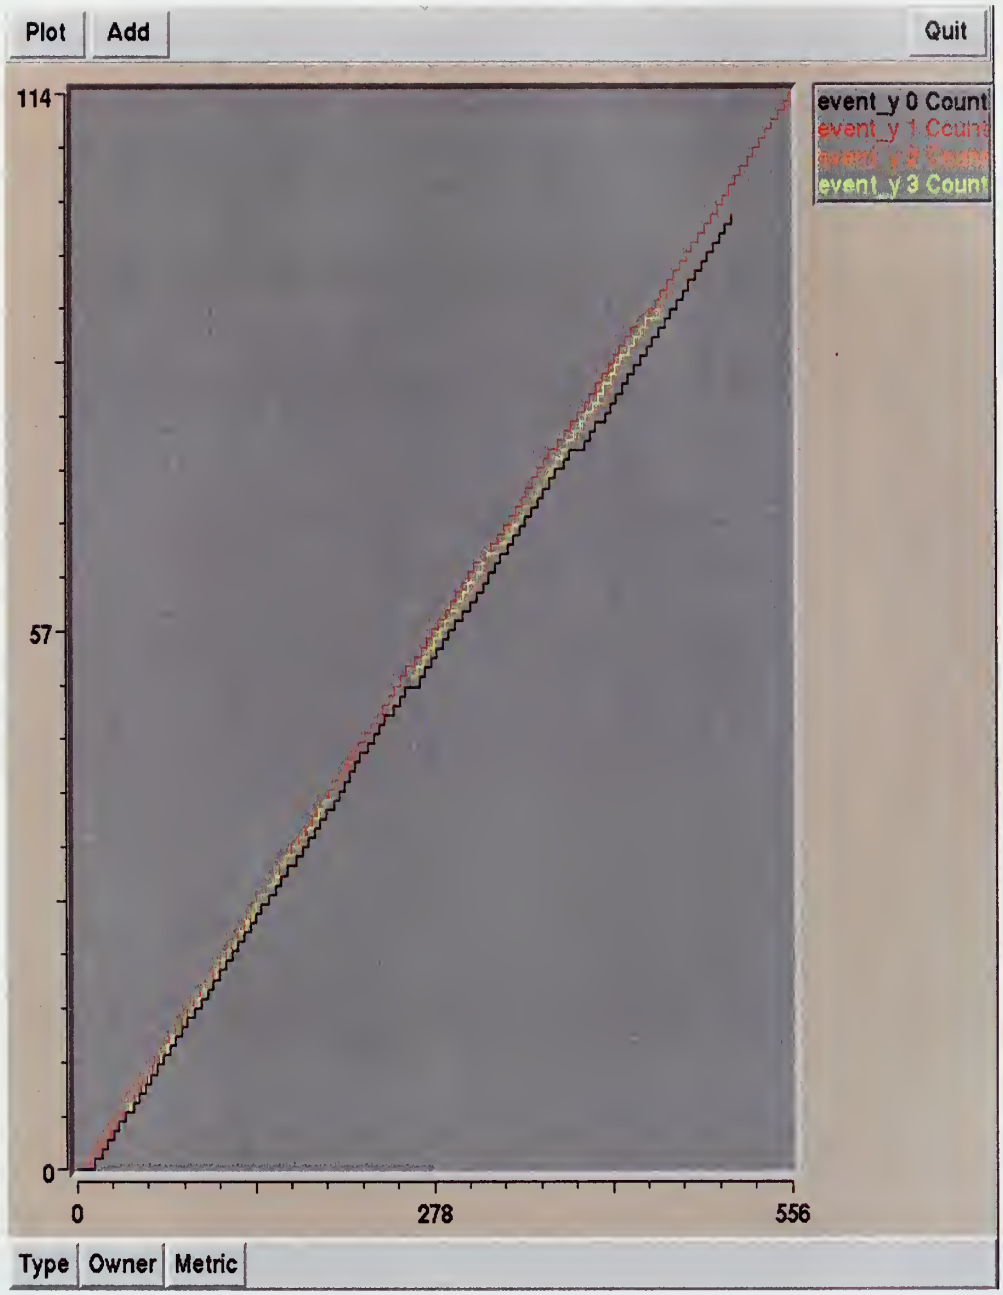

Figure 25. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the GSS scheduling under a low system load condition. Each event_y denotes the completion of a single loop iteration. The Time increases to the right and the number of counts is indicated on the vertical axis.

Figure 26. Plots of event_y for 4 lightweight processes using GSS scheduling under a low system load condition (from the same data as in Figure 25). Each event_y denotes the completion of a single loop iteration. The Time increases to the right.

Figure 27. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the SIMPLE schedule type under a high system load condition. Each event_y denotes the completion of a single loop iteration. The execution time of this case is significantly longer than the SIMPLE schedule, low system load case (see Figure 17). Three of the processes have finished their loop iterations, this is indicated by vertical lines on the left edge of the graph that reach the count value of 100. The remaining process did not complete its work until much later. The Time increases to the right and the number of counts is indicated on the vertical axis.

Figure 28. Plots of event_y for 4 lightweight processes using SIMPLE scheduling under a high system load condition (from the same data as in Figure 27). Each event_y denotes the completion of a single loop iteration. The Time increases to the right. The clusters of events on the left edge of the graph show that Processes 0, 1, and 3 have finished their portion of loop iterations, where as the process 2 did not complete its portion of loop until much later, as shown by the cluster of events on right edge of the graph.

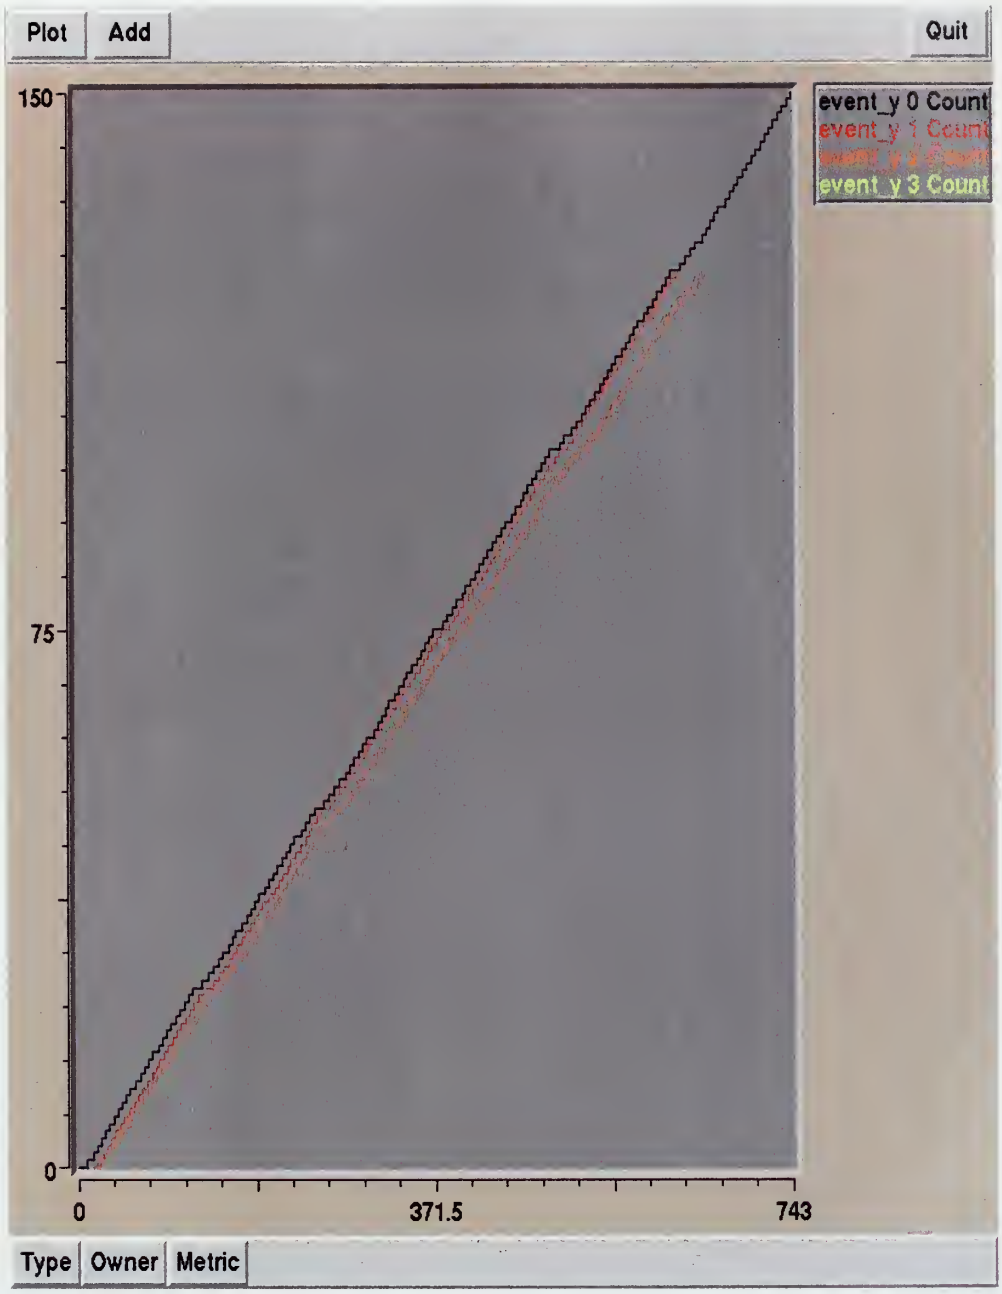

Figure 29. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using DYNAMIC scheduling under a high system load condition. The Time increases to the right and the number of counts is indicated on the vertical axis. Although the workload is not evenly distributed among the processes, even under the high load condition, they finish about the same time.

Figure 30. Plots of event_y for 4 lightweight processes using DYNAMIC scheduling under a high system load condition (from the same data as in Figure 29). Each event_y denotes the completion of a single loop iteration. Although we request 4 processes, only 3 processes participated in the parallel region due to high system load.

program running with the GSS scheduling policy under high load condition does not degrade significantly. This is not the case with the example program running with the INTERLEAVE schedule type under the high load condition (Figure 33 and 34). Like the SIMPLE schedule type under the high load condition, the example program running with the INTERLEAVE scheduling policy can experience a significant slow down if one of the participating processes is blocked. The remaining processes must wait until the last process finishes. As shown in Figure 34, process 0, 1, and 2 spend a large portion of time idling.

The results from running the example FORTRAN program with different workload schedules indicates that the DYNAMIC and GSS scheduling type have better overall performance characteristics as the system load increases, at least when the computation is simple. The strategy of allowing processes to grab more chunks of workload (i.e., be more greedy) as they become available for work is a better overall strategy. It minimizes the amount of wait time when the system load is high and other processes are competing for the available CPU resources.

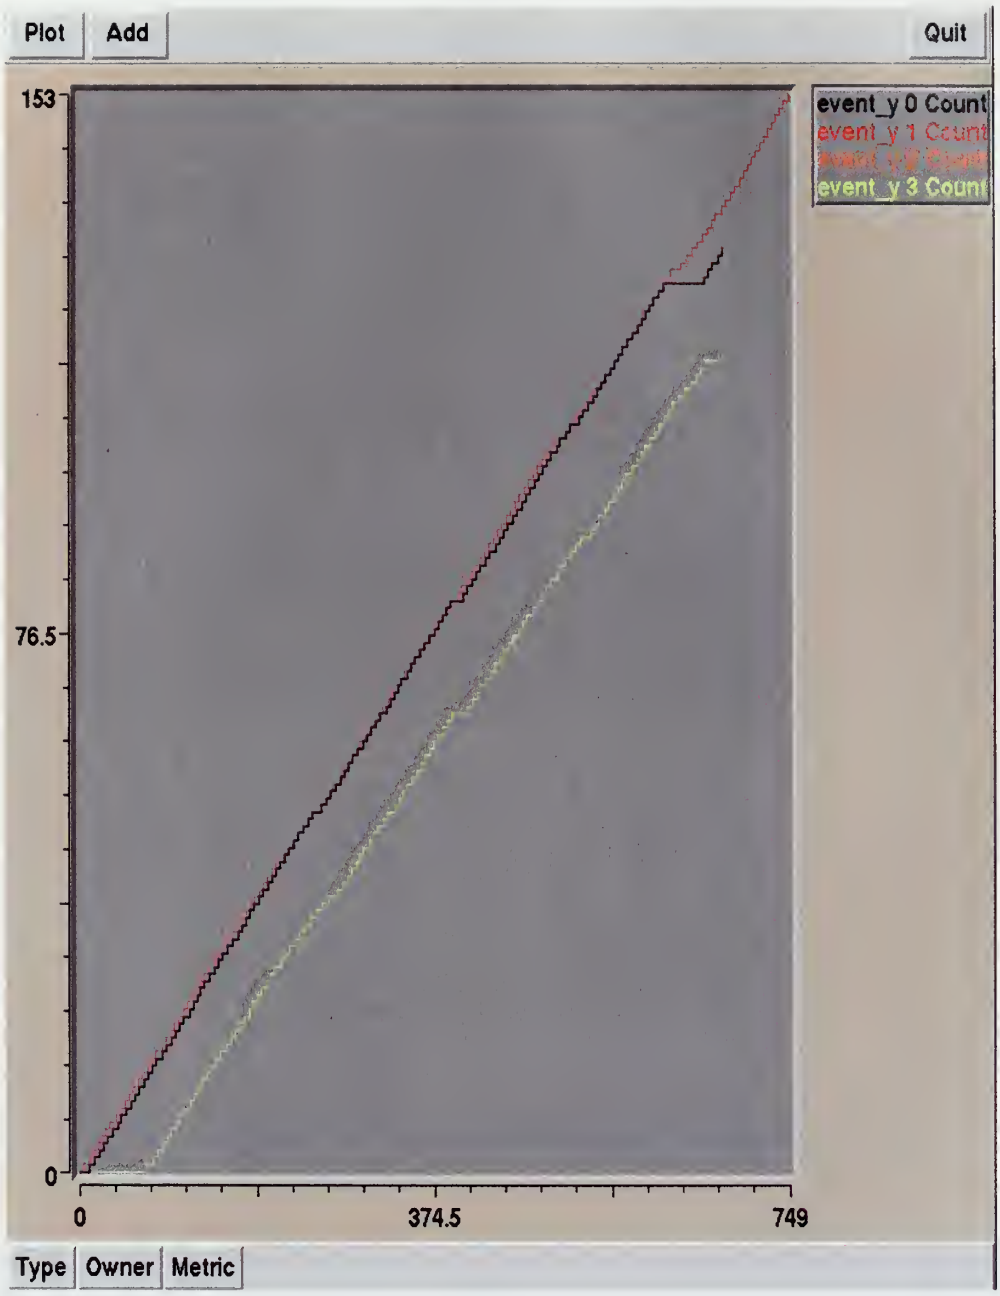

Figure 31. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the GSS schedule type under a high system load condition. The Time increases to the right and the number of counts is indicated on the vertical axis. Although we request 4 processes, only 3 processes actually participated in the parallel region. The workload is not evenly distributed among the processes but they finish about the same time.

Figure 32. Plots of event_y for 4 lightweight processes using GSS scheduling under a high system load condition (from the same data as in Figure 31). Each event_y denotes the completion of a single loop iteration. Although we request 4 processes, only 3 processes (0, 1, and 3) actually participated in the parallel region due to high system load.

Figure 33. Plots of the counts of event_y as a function of time. The plots are for 4 lightweight processes using the INTERLEAVE schedule type under a high system load condition. Each event_y denotes the completion of a single loop iteration. The execution pattern of this case is similar to SIMPLE schedule under a high system load condition. (see Figure 27). Three of the processes have finished their loop iterations, this is indicated by vertical lines on the left edge of the graph that reach the count value of 100. The remaining process did not complete it work until much later. The Time increases to the right and the number of counts is indicated on the vertical axis.

Figure 34. Plots of event_y for 4 lightweight processes using INTERLEAVE scheduling under a high system load condition (from the same data as in Figure 33). Each event_y denotes the completion of a single loop iteration. The Time increases to the right. The clusters of events on the left edge of the graph show that Processes 0, 1, and 2 have finished their portion of loop iterations, where as the process 3 did not complete its portion of loop until much later, as shown by the cluster of events on right edge of the graph.

4. Summary

In this chapter, we discussed the changes we had to make to *Glimpse* for it to monitor automatically parallized FORTRAN programs. In particular, we show that the language independent concepts (i.e., event and interval) and the Data Collection Facility of *Glimpse* is easily adapted to monitor automatically parallelized FORTRAN program.

Using the modified *Glimpse*, we evaluate the performance of a parallel example FORTRAN program with different workload scheduling algorithms under different load conditions. The results indicate that dynamic scheduling schemes (i.e., DYNAMIC and GSS schedule), although having higher overhead due to workload distribution, can perform better when the system load is high. Other scheduling policies performance can degrade significantly under the high system load condition when one or more participating processes is blocked due to competition with other processes.

VII. SUMMARY

In this thesis we describe the design, implementation, and testing of a language independent performance debugging tool called *Glimpse*. The *Glimpse* toolkit collects event data from the execution of multi-threaded programs and provides visualization tools to help programmers analyze the collected data. This work builds upon the approach taken by *Graze*, a performance debugging tool that monitors C++ programs that use the Solaris thread library, and extends it to applications using Java threads and to automatically parallelized FORTRAN programs.

1. Future Work

Before summarizing our experiences with *Glimpse*, we describe some future work areas that will enhance *Glimpse*.

a. Adding generalized math expressions to the interval definition

In chapter IV, we describe the generalization of the *interval* definition to allow for partial matching of attribute values and Boolean relationships between the event's attributes. An example of this interval definition is:

```
interval Transit[ s:Sendmsg -> r:Recvmsg ]
        { s.src=r.src && s.msgsize = r.msgsize } = line;
```

In this example, interval `Transit` requires the attribute value `src` of Sendmag and Recvmsg to match, and that the `msgsize` value of the `Sendmsg` event be equal to the `msgsize` value of the `Recvmsg` event. It would be useful to add generalized mathematical expressions to the interval syntax to further generalize the interval definition. This capability would allow the user to define intervals by specifying certain threshold values. For example, to define intervals where the length of message sent is greater than 100 words, we would write

```
interval Transit[ s:Sendmsg -> r:Recvmsg ]
        { s.src=r.src && s.msgsize/WORDSIZE > 100 } = line;
```

b. *Enhancement to the visualization tools*

Currently, the interval definition is defined in the specification file. This file is read by the visualization tools, *gorge* and *nibble*, during start up. To visualize the event data with a different interval definition, the user needs to modify the specification file and restart the visualization tools in order for it to display the changes to the interval definition. It would be useful to add the capability, in the visualization tools, to allow the user to create new interval definitions, and to edit the existing interval definitions on the fly. This would make it less cumbersome for the user to experiment with different interval specifications.

Another enhancement to the visualization tools is to add the capability to associate the log data being displayed with the thread name from which it came. Recall from the discussion in chapter IV, each thread is mapped to a unique log file by the Data Collection Facility. In the current implementation, the logical thread name used by the program during execution is not saved to the log file. Therefore, to associate data from a log file back to a particular thread requires some careful interpretation of plots and understanding of the program code by the user. To make this process easier for users, we should modify the Data Collection Facility to save the logical name of the thread along with the data to the log file. The visualization programs can then use this information to label the plots with the thread names.

c. *Improving the performance of Glimpse's visualization tools for large data sets*

After we successfully tested the FORTRAN version of *Glimpse* with various workload scheduling policies (chapter VI), we applied it to a Navy Operational Weather Forecast application. This application, COAMPS, is a weather simulation program written in FORTRAN. The current operational version of COAMPS relies heavily on the compiler-directive approach to optimize the runtime performance. However, we noticed that its performance does not scale well in the multiprocessor environment. For example, the maximum speedup of this application when running

on SGI multiprocessor systems is approximately 2.6 times that of the single processor mode. This maximum speedup value is observed when the COAMPS application is running with 4 processors.

We instrumented this application with *Glimpse* in an attempt to identify the possible performance bottlenecks. The instrumented code did not add any significant overhead to the COAMPS run time; the instrumented code is perhaps 2 to 3 percent slower than the non- instrumented code. The data volume generated by the instrumented code was about 50 mbytes (approximately 250,000 events) for a typical COAMPS run. The *gorge* tool took almost an hour to load the event data and display the intervals. The slowness appears to be related to the algorithm that is used to build the interval instances from the event data stream. Given a number of event instances in the event stream n, this algorithm finds all instances of an interval type by starting the search at the head of the event stream and moving forward until it finds a matching ending event. This algorithm performs this operation n times, each time starting the search one position forward in the event stream. Thus, for the worst case scenario, the total number of operations needed to find all instances of a single event type is given by $(n\text{-}1) + (n\text{-}2) + (n\text{-}3) + \ldots + 2 + 1$. This summation reduces to $n * (n - 1) / 2$, which is on the order of n^2 for large n. Further investigation is needed to improve the performance of this algorithm for handling large event data sets.

2. Experiences using *Glimpse* with concurrent Java and FORTRAN programs

We found that the `event` and `interval` concepts, together with visualization tools that can display them, are very useful for programmers. It helps them to understand the behavior of parallel applications, independent of the language environment. The glimpse's *event* and *interval* concepts are similar to that of the *event* and *state* used in the commercial software, *TimeScan* (see chapter II). However, in TimeScan the event definitions are defined in the source code, whereas *Glimpse* uses a speci-

fication file to store the user-defined events. The specification file approach is more flexible and minimizes the changes to the source code of the program that the user wants to monitor. Another important difference between *TimeScan* and *Glimpse* is that *TimeScan* does not support an interval that spans two threads. *Glimpse* does not have this limitation.

Naturally, the process of porting *Glimpse* to the Java and FORTRAN environment, and to other languages or programming environments, requires writing a platform-specific Data Collection Facility that contains such utilities as functions to handle input/output, timestamps, a utility program to translate event specifications to data logging functions, and to handle other low level issues. From the perspective of the user of the *Glimpse* tool, even such implementation differences are hidden. The user only needs to concentrate on understanding the behavior of the program, and describing those behaviors using events and intervals.

APPENDIX A. OVERVIEW OF JAVA AND JAVA THREAD

This appendix presents a brief overview of Java and the mechanisms that it provides for building multi-threaded programs.

a. Overview of Java

Java is an object-oriented programming language developed by Sun Microsystems [Ref. 20]. One of its design goals is to allow programs to run on different type of operating systems and architectures without recompilation or any other modification. To achieve this goal of platform independence, the Java specification describes not only the programming language itself, but also the Java Virtual Machine (JVM) [Ref. 21] and Application Programming Interface (API). The concept of the Java Virtual Machine is similar to that of a virtual machine operating system. It defines the interface and functionality of the layer between Java programs and that of the underlying architecture. A Java program is compiled into bytecodes, which are instructions understood by the JVM, and stored in class files. The class files are then loaded and executed by the Java Virtual Machine. To execute a Java program, a Java Virtual Machine needs to provide support for (i) memory allocation, which includes garbage collection, (ii) dynamic loading and verification of class files, (iii) index checking on array access, and (iv) threads. It also provides an interface to native functions such as those that support network and windowing environment operations. Currently most implementations of the Java Virtual Machine are implemented in software and thus depend on resources provided by the underlying native operating system. Consequently there is some overhead and performance penalties associated with executing bytecodes compared with executing typical native machine code.

In addition to the Java Virtual Machine, the Java specification defines a comprehensive API specification. In the Java version 1.1, there are approximately 500 classes for the core library alone, and this number is expected to get larger in the

next release. The core Java API, in essence, defines the software infrastructure that should be available on all platforms that support Java. The current API includes packages for graphical user interface components, thread libraries, input and output streams, data structures, database access and internationalization of applications.

By providing a high-level and uniform API across various platforms, the Java environment isolates programmers from the variations and differences in functionality and interfaces provided by different operating systems. For example, Java provides a Thread class (java.lang.Thread) that contains constructors for creating a thread and methods for setting the priority of a thread, suspending a thread, naming a thread and getting information about a thread's status. This package has the same interface on both POSIX and Window NT platforms. This standardization of interface permits programmers to focus on building programs rather than on learning another programming interface. It is the job of the commercial vendors and the freeware and open-source community to provide an efficient and correct implementation of the Java API library and JVM for their respective platforms.

In section (b.), we will describe the synchronization mechanism that is built into the Java programming language, specifically, the monitor object, and how it is used to ensure that a class method can have exclusive access to a variable or another object in a Java program. Additionally, the wait and notify mechanism provided by Java to transfer execution between threads is discussed. In the conclusion section, we will summarize a few simple rules for using synchronization in a Java program.

b. Java Threads and Synchronization

The Java language is designed to support concurrent programming, with the thread as the basic unit of concurrency. A concurrent Java program can have many threads of execution. In some implementations of the Java Virtual Machine, threads can execute on a separate processor if the underlying hardware is a multiple-processor system. In such systems the execution of threads can overlap in time. On single-processor systems, the Java Virtual Machine supports concurrent execution by inter-

leaving different threads to give the appearance of simultaneous execution.

A fundamental requirement for support of concurrent programming in any language is the ability to provide synchronized access to a shared variable or to any of a set of related critical sections. In many concurrent programs mutual exclusion is the technique used to ensure proper synchronization among threads. The following discussion will focus on the language features provided by the Java programming language for implementing synchronization using mutual exclusion.

A typical multi-threaded Java program consists of cooperating threads working on some tasks, where each thread has its own program counter and stack. Within the context of each execution thread, program instructions are carried out in sequential order. Variables that are used, or assigned to, while each thread executes can be considered as either residing in local or shared memory. Each thread has its own copy of local variables, which are not visible to any other threads. A shared variable, on the other hand, is visible to any thread that has a reference (i.e., a memory address) to that variable. Cooperating threads can only pass information to one another through shared variables. Therefore, the issue of synchronization in general, and specifically mutual exclusion, is only of concern for accessing variable in shared memory locations.

The simple Java program below can be used to illustrate that simultaneous access to shared variables, without proper synchronization, could give an incorrect result. The program models the transaction between customers and their bank accounts.

```
class bank_account {

    /* The current dollar amount in the customer account */
    int balance;

    /* constructor, initialize account. */
    bank_account(int initial_deposit) {
        balance=initial_deposit;
    }
```

```
    void Deposit(int deposit) {
        balance=balance+deposit;
    }

    void Withdraw(int amount_to_withdraw) {
       /* check for sufficent funds in account */
       if (balance >= amount_to_withdraw)
       {
          balance = balance-amount_to_withdraw;
       }
       else
       {
          /* issue warning about account overdrawn */
          System.out.println("Account overdraft");
          return;
       }
    }
}
```

The **bank_account** class has two methods: the **Withdraw** method that allows the customer to remove money from the account and the **Deposit** method that lets the customer put money in. The account is initially credited with some dollar amount when it is first created. To ensure overdraft does not occur, the **Withdraw** method first verifies that the account has sufficient funds to cover the transaction. If there are insufficient funds, a warning message is issued and the **Withdraw** method returns without changing the balance. Otherwise the balance of the account is reduced by the amount withdrawn.

To simulate customer activity on the bank account, we let each customer's transaction be represented by a thread. Suppose two customers try to withdraw money from the same account. The **Withdraw** method of the **bank_account** object is invoked twice, once by each separate threads. Each thread uses a different copy of the variable **amount_to_withdraw**, but modifies the same copy of the variable **balance**, i.e., the variable **amount_to_withdraw** resides in the local memory of each thread whereas the variable **balance** resides in the shared memory. In summary, more

than one thread can be executing in the same instruction code segment, but they each have their own program counters and local working memory, and can access or modify variables shared with other threads.

Before describing why synchronization is important, we need to look at how a JVM schedules threads for execution. In the context of scheduling, a Java thread that is not currently executing is either in a *runnable* or *blocked* state. Typically, the *runnable* state is the default state for a thread unless it has performed an operation that causes it to become blocked. For example, a thread that performs a network-related operation such as reading from a socket when the data for that socket is not yet available will become blocked. When the data becomes available the thread will become runnable again. In a multi-threaded Java program, every thread is competing for CPU cycles and it is up to the JVM to select a thread from among the pool of *runnable* threads to be the currently running thread. JVM uses a priority-based preemptive scheduling mechanism that guarantees the currently running thread always has a priority value equal to the highest priority value found among all *runnable* threads.

A thread will remain as the currently running thread until a scheduling event signals the JVM to select another thread from the pool of *runnable* threads as the current running thread. There are three types of scheduling events: (1) when a currently running thread becomes blocked or completes; (2) when a higher priority thread enters the *runnable* state; and (3) when a timer interrupts execution. In the case of scheduling event (1) or (2), the currently running thread yields the CPU and the JVM selects another runnable thread. Not all implementations of a JVM must support the scheduling event (3). On most UNIX platforms only the first two scheduling events are supported, whereas on the Windows NT operating system, the timer-based scheduling event is also supported. A multi-threaded Java program will behave differently depending on the JVM scheduling events. For example, consider a Java program that has two long running threads of equal priority, and assume that

neither thread performs an instruction that will cause it to become blocked. On UNIX systems, if thread 1 is started first and never enters the blocked state, it will continue to run until it is completed. When thread 1 is completed, it leaves the *runnable* state, thus creating a scheduling event, at which time thread 2 becomes the currently running thread. Running the same program on a JVM that supports timer-based scheduling, thread 1 starts first and at some time later a timer-based scheduling event occurs. Thread 2 becomes the newly running thread and is executed by the CPU until the next timer-based scheduling event, at which time thread 1 becomes the running thread again. Under the timer-base scheduling scheme, thread 1 and 2 alternate as the currently running thread (i.e., similar to a round-robin scheduler), with each thread getting a time-slice of CPU cycles. Thread 2 might even finish before thread 1.

From the perspective of a running thread, it does not know when a scheduling event will occur. The switching of execution context from one thread to another thread is performed by the JVM; a running thread can be interrupted at any time as the result of any of the scheduling events described above. We now return to the synchronization issue and our example program. We will assume the **bank_account** program is running on a JVM with timer-based scheduling.

In the following scenario, suppose the bank account is a joint account and that two people attempt to obtain cash from their joint account at the same time. For example, the husband (thread A) is withdrawing $80 and the wife (thread B) is withdrawing $50 from the same account. Assume the initial balance is $100. If after thread A executed the first line of code

```
(if (balance >= amount_to_withdarw)),
```

but before it can execute the next line of code to subtract the amount from the variable **balance**, a timer-based scheduling event occurs. The JVM will now select the next thread from the *runnable* thread pool with the highest priority value as the current running thread. Since both thread A and B have the same priority, the JVM

will move thread A to the *runnable* pool and make thread B the currently running thread. Thread B executes the following line of code

```
(if (balance >= amount_to_withdarw)).
```

The variable **balance** is still at $100 because thread A was interrupted before it could modify the value of **balance**. Thread B continues to execute and subtracts $50 from the joint account, and the variable **balance** is updated to reflect the new value of $50. Thread B finishes execution by exiting the **Withdraw** method and a scheduling event is generated. The JVM selects thread A as the new running thread. Thread A resumes execution at the following line of code:

```
balance=balance-amount_to_withdraw;
```

The variable **balance** has the value of $50 at this point, and now thread A subtracts $80 from $50. The new value of the variable **balance** is negative $30 dollars, which is clearly incorrect. This error results from lack of synchronization when multiple threads try to access and modify a shared variable. In this example, a mechanism is needed to provide proper synchronization to the shared variable **balance**. Such a mechanism would immediately block thread B when it tries to execute the **Withdraw** method, because thread A was already in this same method. This type of synchronization is called mutual exclusion. Mutual exclusion ensures that only one thread can enter a particular code segment or method that modifies and uses the same shared variable. These code segments or methods are called critical sections. If mutual exclusion is used in this example, when thread B tries to execute the **Withdraw** method, it will be blocked and the JVM will re-select thread A for execution. After thread A completes, thread B is unblocked and is made the currently executing thread by the JVM. Thread B now sees variable **balance** has a value of $20 and realizes that there are insufficient funds for withdraw.

Java supports mutual exclusion with a mechanism known as a monitor. The concept of the monitor was first introduced by Brinch Hansen [Ref. 22, 23]. A

monitor acts as a lock to a section of program code, allowing only a single thread to enter that code segment at a time. Any object in Java can be a monitor. To place a block of code under the auspices of a single monitor, the Java language provides a reserved keyword *synchronized.* For example, to ensure that variable **balance** is properly updated when accessed by different threads, we modify the **bank_class** by adding the *synchronized* statement to the block of code that encloses the variable **balance**.

```
void Withdraw(int amount_to_withdraw)
{
   synchronized(this) { /* Acquire monitor, beginning of the
                           critical section, use 'this' object
                           as the monitor */

       if (balance >= amount_to_withdraw) /* check for sufficent
                                             funds in account */
       {
           balance = balance - amount_to_withdraw;
       }
       else
       {
           /* issue warning about account overdrawn.*/
           System.out.println("Account overdraft");
           return;
       }
   } /* Release monitor, end of the critical section */
}
```

The statement *'synchronized (this)'* marks the beginning of the critical section to be placed under the monitor. The token following the 'synchronized' is the reference to the object acting as the monitor for this code segment. In this example *'this'* refers to the object which the **Withdraw** method belonged to. Other objects can also be used as the monitor in place of *'this.'* For instance, we could have defined a string object in the class declaration of **bank_account** and used that string object as the monitor and it would have the same effect.

```
class bank_account {

    /* Monitor object for each instance of the bank_account,
       each account has its own monitor. */
    String customer;

    /* The current dollar amount in the customer account */
    int balance;

    /* constructor, initializes the customer account. */
    void bank_account(String customer_name, int initial_deposit)
    {
        customer=new String(customer_name);
        balance=initial_deposit;
    }
    .
    .
    void Withdraw(int amount_to_withdraw)
    {
        synchronized(customer) /* Acquire monitor at the
                                  beginning of the critical
                                  section, use the string
                                  variable 'customer'
                                  as the monitor. */
        {
           if (balance >= amount_to_withdraw) /* check for sufficient
                                                 funds in account */
           {
               balance = balance - amount_to_withdraw;
           }
           else
           {
               /* issue warning about account overdrawn */
               System.out.println("Account overdraft");
               return;
           }
        } /* Release monitor, end of the critical section */
   }

}
```

Before the current thread (the thread that the CPU is currently executing

instructions for) can execute code inside the critical section, it must first enter the monitor guarding that critical section. If another thread has already entered any critical section protected by the same monitor, then the current thread must wait until the other thread exits the monitor. Many critical sections can be associated with the same monitor, but only one thread can be granted that monitor at any one time. In other words, if one thread is executing in a critical section guarded by a monitor, than no other threads can enter other critical sections guarded by the same monitor. Conversely a single thread can enter more than one monitor if need be. For instance, a thread can execute in one critical section guarded by monitor A and then enter another critical section guarded by monitor B, provided that no other threads have entered monitor B. This is perfectly legal as the mutual exclusion mechanism in Java is based on scope of the monitor object.

If a thread has been granted the ownership of a monitor, and tries to acquire the same monitor again the thread would not be blocked. This situation arises when the currently executing thread, while inside a critical code section tries to execute another code segment guarded by the same monitor. For example, a thread making recursive calls from a synchronized method. The design of Java monitors allows the same thread to re-enter the critical section once it acquires the monitor, or to enter other critical sections guarded by the same monitor. This re-entrant behavior of the monitor eliminates the possibility of a thread waiting on itself.

Returning to the bank account example, we consider the situation that occurs when one customer tries to withdraw more money than is available in the account. Instead of returning with an error message, we may prefer that the customer wait until another customer has deposited additional money into that account to cover the withdrawal. This is a very common pattern in many multi-threaded programs where a thread executing in a critical section might encounter a condition such that it cannot proceed until another thread can act to change that condition. What is needed to support this type of programming construct is a mechanism to transfer

the execution from one thread to another thread in the critical section. In order for this to work, the currently executing thread must temporarily halt and give up its claim to the monitor so that another thread can execute in that critical section. Once the condition is satisfied by the action of another thread, the original thread can then proceed again. The Java language provides a wait and notify mechanism for this purpose. Of course, the inclusion of wait and notify are necessary to make the mechanism that we have been discussing truly qualify as a monitor. Without them, semaphores may have been included in the language and would likely have been more efficient.

To see how wait and notify work, we modify our bank account program so that when a customer overdraws an account, he does not receive an error message but instead merely waits until another customer deposits more money into that account.

```
class bank_account {

    /* Monitor object for each instance of the bank_account,
       each account has its own monitor. */
    String customer;

    /* The current dollar amount in the customer account */
    int balance;
    .
    .
    void Deposit(int deposit)
    {
        synchronized (customer)
        {
           balance=balance+deposit;
           customer.notifyAll();    /* Notify any thread waiting on
                                       this monitor */
        }
    }

    void Withdraw(int amount_to_withdraw)
    {
       synchronized (customer)        /* Acquire monitor */
       {
```

```
            while ( true )
            {
               if (balance >= amount_to_withdraw) /* check for
                                                     sufficient fund */
               {
                  balance = balance - amount_to_withdraw;
                  return;
         }
               else
               {  /* Relinquish monitor and put the current
                     thread on the wait queue of this monitor.
                     Catch interrupted exception if this thread
                     unexpectedly returns from wait() method. */
                  try { customer.wait(); }
                  catch (InterruptedException e) { System.exit(1); }
               }
            }
         } /* End of critical section */
      }
}
```

The modified code works as follows. If a thread tries to withdraw more money than what is in the specified bank account, it will be suspended from execution (i.e., it is blocked in the **wait()** method call) until some other thread has called the **Deposit()** method. After another thread deposits money into the account, it calls the **notifyAll()** method before exiting the **Deposit()** method. The **notifyAll()** method sends a signal to all threads that are blocked in the **wait()** method. These threads are removed from a wait queue and placed in the ready queue by the JVM. When a blocked thread returns from the **wait()** method it checks whether there are sufficient funds in the account. If not, it calls the **wait()** method again and repeats the process until other threads have deposited sufficient money to cover the withdrawal. The **wait()** and **notifyAll()** can only be called from inside a monitor, i.e., only the thread currently inside the monitor can invoke the **wait()** and **notifyAll()** methods of that monitor.

Unlike the **wait()** method, the **sleep()** or the **yield()** method (from the

java.lang.Thread library) has no effect on the state of the monitor when invoked. It merely tells the JVM to select another thread as the currently running thread. When the **wait()** method is invoked, it places the current thread on the wait queue of the monitor and relinquishes mutual exclusion prior to suspending the current thread. The **notifyAll()** method releases all threads waiting on a condition in that same monitor. The **wait()** and **notifyAll()** methods are not part of **java.lang.Thread** library but are an integral part of the Object class and are tightly integrated into the JVM. They provide a reliable mechanism to safely transfer the control from one thread to another thread in a monitor.

Suppose instead of calling the wait() method, a thread invokes the **sleep()** or **yield()** method while inside the **Withdraw()** method to temporarily suspend itself so that another thread can be executed by the JVM. In this example, calling the **sleep()** or **yield()** method would lead to a dead-lock situation when executing on a machine with a non-preemptive thread scheduler: one customer starts to perform a withdraw action but cannot complete that transaction until more money is deposited by another customer. However the deposit action by another customer cannot begin until the **Withdraw** method is completed because the depositing thread would not be permitted to enter the monitor.

The Java language provides the basic synchronization primitives via the monitor and the wait and notify mechanisms. Other, higher-level, synchronization facilities can be built on top of these basic primitives. For example, counting semaphores, can be built from the **wait()** and **notifyAll()**. In the basic wait and notify mechanism if a **notifyAll()** method is called when there are no threads waiting on a monitor, then that notification event has no effect. Counting semaphores, however, maintain state information that indicates how many times **wait()** or **notifyAll()** are called. When one thread calls the **notifyAll()** method and there are no threads waiting, the notification event causes a value to be incremented in the data structure of the associated semaphores. Counting semaphores are useful for providing synchroniza-

tion to a shared buffer in a producer-consumer scenario. Similarly, other higher-level synchronization facilities such as those needed to solve the multiple reader, single writer problem can be constructed using the basic primitives provided by Java.

c. Summary of Java Threads and Synchronization

In this appendix, we reviewed the support for writing multi-threaded program provided by Java and some basic concepts for synchronization. These include:

a. Cooperating threads that communicate via shared variables. If the correctness of program execution depends on the order in which threads act on the shared variable, then access to these variables needs to be synchronized.

b. Java provides an elegant and simple mechanism for synchronization. To ensure synchronized access to shared variables, code segments that act on shared variables should be placed inside a synchronized block or method.

c. Before a thread can execute a synchronized block or method, it first acquires the monitor associated with that block or method. A monitor can be associated with multiple synchronized blocks or methods. However, only one thread can execute inside a synchronized block or method of a given monitor at any time.

d. To transfer control of program execution between threads from inside a synchronized block or method, a program should use **wait()** and **notifyAll()** methods provided in the monitor class. These methods guarantee that the monitor is relinquished when a thread is placed in a wait state, and that the monitor is re-entered prior to resuming execution.

e. Other forms of synchronization techniques such as the counting semaphore for controlling multiple access to limited resources, and mechanisms that can be used to solve such commonly occurring problems as the multiple-reader, single-writer can also be constructed using the basic monitor mechanism provided by Java.

APPENDIX B. GRAMMAR OF THE GLIMPSE SPECIFICATION LANGUAGE

This appendix describes the grammar for the *Glimpse* specification language. The italicized words are non-terminals; words enclosed by single quotes are literals. Tokens that denote keywords used by the language are prefixed with k_. The list of keywords is given in section two.

a. Grammar of the Glimpse Specification Language

The goal of specification language:

spec ::= k_glimpse *ident events intervals* k_end '.'

`Event` has a name, associated attributes, and an optional graphic symbol.

events ::= NULL | *event events*
event ::= k_event *ident memberlist* '=' *symbol* ';' | k_event *ident memberlist* ';'

`Event` can have 0 or more attributes; attributes are separated by commas and are enclosed by a left and right parentheses.

memberlist ::= NULL | '(' *members* ')'
members ::= NULL | *ident* | *ident* ',' *members*
intervals ::= NULL | *interval intervals*

An `interval` has a name, a start event, search direction, an ending event, a matching criteria, and a graphic symbol.

interval ::= k_interval *ident* '[' *ident search_dir ident* ']' *match_type* '=' *connection_type* ';'

A second form of `interval` production is presented below. This production allows a boolean expression for matching the start and end events.

interval ::= k_interval *ident* '[' *ident* ':' *ident search_dir ident* ':' *ident* ']' *where_clause* '=' *connection_type* ';'

search_dir ::= ',' | k_forwards | k_backwards

The production below specifies the matching criteria for the start and end event. If the match_type is NULL, then matching is done based on the start and end event type matching only these events that occur within the same thread. Otherwise, the matching is done on the attribute values.

match_type ::= NULL | k_match

The logical expression for matching the start and end events is shown below.

where_clause ::= NULL | '' *expr* ''
expr ::= *relation* | *relation* k_and *expr*
relation ::= *ident* '.' *ident math_oper ident* '.' *ident*
math_oper ::= k_lt | k_le | k_eq | k_ge | k_gt | k_ne
connection ::= k_connect '(' *connection_type* ')'

Graphics symbols for `event` and `interval` types are shown below.

symbol ::= *symbol_type offset* | k_symbol '(' *symbol_type* offset ')'

symbol_type ::= NULL | k_x | k_plus | k_box | k_diamond

connection_type ::= NULL | k_line | k_rectangle

offset ::= NULL | ':' ident

ident ::= k_ident

b. Keywords for the Glimpse Specification Language

The keywords used in the *Glimpse* specification language are on the right hand side of the statements below, shown as regular expression. As in the specifications for `lex`, [a-zA-Z] means any character between `a` and `z` or between `A` and `Z`. The * represents the Kleene closure.

k_graze ::= `graze | Graze`

k_event ::= `event | Event`

k_end ::= `end | End`

k_symbol ::= `symbol | Symbol`

k_connect ::= `connection | Connection`

k_interval ::= `interval | Interval`

k_ident ::= `[a-zA-Z][a-zA-Z0-9_]*`

k_plus ::= `plus | Plus`

k_x ::= `x | X`

k_box ::= `box | Box`

k_diamond ::= `diamond | Diamond`

k_rectangle ::= `rectangle | Rectangle`

k_line ::= `line | Line`

k_where ::= `where | Where`

k_match ::= `match | Match`

k_lt ::= `<`

k_le ::= `<=`

k_eq ::= `==`

k_ge ::= >=

k_gt ::= >

k_ne ::= !=

k_and ::= &&

k_or ::= ||

k_not ::= !

k_number ::= [0-9]+

k_backwards ::= <-

k_forwards ::= ->

APPENDIX C. A NOTE ON EXPERIENCE USING THE FINALIZE FACILITY OF JAVA'S GARBAGE COLLECTION

This appendix discusses the experience we had using the Java Garbage Collector to automatically invoke the clean-up methods needed by the *Glimpse* Data Collection Facility.

The Java programming language provides a `finalize()` method that is invoked by the JVM for an object that is about to be garbage collected. The `finalize` mechanism allows an object to perform any clean up actions such as closing the file that it currently has open, before that object is freed by the system. We experimented with putting the code to properly close a log file inside a `finalize` method of a `ThreadData` object, thus eliminating the need to explicitly call the `closeThreadData()` and `close()` methods of the `GzEvent` object. One problem we encountered is due to the asynchronous nature of the Garbage Collector in the JVM. After a thread has exited its run method, the `finalize()` method might not be called until the execution of the program has continued on for quite some time. We wrote a short program to test the behavior of the Java Garbage Collector:

```
import java.util.*;
import java.io.*;
class invoke_final extends Thread{
    public invoke_final()
    {
        super();
        System.out.println("Create object: "+this.getName());
    }
    protected void finalize() throws Throwable {
        try {
           System.out.println("finalize method called on id "+
                              this.getName());
        } finally {
           super.finalize();
```

```
        }
    }
    public void run()
    {
        System.out.println("running id: "+this.getName());
    }
    public static void main(String[] args) {
        /* Forever loop that creates new threads. */
        while (true) {
            /* Create some threads and run them. */
            for (int i=0; i<100; i++) {
                invoke_final x=new invoke_final();
                x.start();
                Thread.yield();
            }
            /* Request JVM to run Garbage Collector */
            Runtime.getRuntime().gc();
            /* Create more threads and run them. */
            for (int i=51; i<100; i++) {
                invoke_final x=new invoke_final();
                x.start();
            }
        }
   }
}
```

In this test program, a forever `while` loop creates threaded objects. Each time through the `for` loop, an object is created and the thread associated with that object is started. At the beginning of each loop, the reference to the previously created object is removed. Since there is no longer any reference to the object in the main program, the JVM sees these unreachable objects as candidates for garbage collection. We inserted a call to the Java Garbage Collector in the middle of the while loop. When the Garbage Collector runs, it should invoke the `finalize()` method, which prints out a string indicating that the finalize method for this object has been invoked by the Garbage Collector.

When we ran this program, we saw different results under different JVM's. On the Windows NT system (SUN Java 1.1.8), at the point where we call

`Runtime.getRuntime().gc()`, we saw outputs indicating that the Garbage Collector ran. On the SGI (SUN Java 1.1.6), the Garbage Collector did not run until the outer `while` loop had gone through 10 to 20 iterations. Each iteration of the `while` loop creates 200 objects.

In summary, we would like the *Glimpse* data collection to close the log file when a given event thread exits, and not to wait until some time later. Unfortunately, the time of execution of the `finalized()` method, under the current JVM specification, is implementation-dependent. For example, placing the clean-up code inside the `finalize()` method works on the Windows NT system, but not on an SGI.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

[1] Lantz Moore, Debra Hensgen, David Charley, Philip A. Wilsey, and Venkatram Krishnaswamy. Graze: A tool for performance visualization and analysis. *International Conference on Parallel Processing*, 1995.

[2] T. Hogan and T. Rosmond. The description of the navy operational global atmospheric prediction system's spectral forecast model. *Monthly Weather Review, 119*, 1991.

[3] R. Hodur. The naval research laboratory coupled ocean/atmosphere mesoscale prediction system (coamps). *Monthly Weather Review, 125*, 1997.

[4] Defense Modeling and Simulation Office. *Master Environmental Library (MEL), Technical Reference Guide, Characteristics and Performance, Version 1.1*, 1998.

[5] T. Hogan and T. Rosmond. The description of the navy operational global atmospheric prediction system's spectral forecast model. *Monthly Weather Review, 119*, 1991.

[6] Sharon E. Perl, William E. Weihl, and Brian Noble. *Continuous Monitoring and Performance Specification.* System Research Center, Report 153, 1998.

[7] Barton P. Miller, Mark D. Callaghan, Jonathan M. Cargille, Jeffrey K. Hollingsworth, R. Bruce Irvin, Karen L. Karavanic, Krishna Kunchithapadam, and Tia Newhall. The paradyn parallel performance measurement tools. *IEEE Computer 28, 11, (November 1995), pages 37-46.*, 1995.

[8] Daniel A. Reed, Ruth A. Aydt, Luiz DeRose, Celso L. Mendes, Randy L. Ribler, Eric Shaffer, Hueseyin Scimitci, Jeffery S. Vetter, Daniel R. Wells, Shannon Whitmore, and Ying Zhang. *Performance Analysis of Parallel Systems: Approaches and Open Problems.* Department of Computer Science, University of Illinois, 1997.

[9] Randy L. Ribler, Jeffery S. Vetter, Huseyin Simitci, and Daniel A. Reed. Autopilot: Adaptive control of distributed applications. *Proceedings of the 7th IEEE Symposium on High-Performance Distributed Computing, Chicago, IL, July 1998.*, 1998.

[10] I. Foster, C. Kesselman, and S. Tuecke. The nexus approach to integrating multithreading and communication. *Journal of Parallel and Distributed Computing, 37*, 1996.

[11] Randy L. Ribler, Huseyin Simitci, and Daniel A. Reed. The autopilot peformance-directed adaptive control system. *Future Generation Computer Systems, special issue (Performance Data Mining), to appear.*, 1999.

[12] Daniel A. Reed, David A. Padua, Ian T. Foster, Dennis B. Gannon, and Barton P. Miller. Delphi: An integrated, language-directed performance prediction, measurement, and analysis environment. *The 9th Symposium on the Frontiers of Massively Parallel Computation, Annapolis, MD*, 1999.

[13] Dolphin Interconnect Solutions, Inc. *TimeScan, Event Analysis System User Guide, version 3.0*, 1997.

[14] Dolphin Interconnect Solutions, Inc. *TimeScan, ELOG Library Programmer's Guide, version 3.0.1*, 1997.

[15] Sharon E. Perl and William E. Weihl. Performance assertion checking. *Proceedings of the Fourteenth ACM Symposium on Operating Systems Principles*, 1993.

[16] Parallel Tools Consortium. *Manual Page for the Portable Timing Routines*, 1996.

[17] Sheng Liang. *The Java Native Interface: Programmer's Guide and Specification.* Addison-Wesley Pub Co., 1999.

[18] Silicon Graphics, Inc. *MIPSpro Fortran 77 Programmer's Guide*, 1999.

[19] Silicon Graphics, Inc. *UNIX MAN Page for MP, Multiprocessing directives for MIPS FORTRAN Compiler*, 1999.

[20] James Gosling, Bill Joy, and Guy Steele. *The Java Language Specification.* Addison-Wesley Longman, Inc., 1996.

[21] Tim Lindholm and Frank Yellin. *The Java Virtual Machine Specification.* Addison-Wesley Longman, Inc., 1996.

[22] P. Brinch Hansen. A comparison of two synchronizing concepts. *Acta Informatica, Volume 1, Fasc. 3, (1972), pages 190-199.*, 1972.

[23] P. Brinch Hansen. Concurrent programming concepts. *Computing Surveys, Volume 5, Number 4, (December 1973), pages 223-245.*, 1973.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center — 2
   8725 John J. Kingman Rd., STE 0944
   Ft. Belvoir, Virginia 22060-6218

2. Dudley Knox Library — 2
   Naval Postgraduate School
   411 Dyer Rd.
   Monterey, CA 93943-5101

3. On-Scene Systems, Naval Research Laboratory, Monterey — 1
   Code 7542
   7 Grace Hopper Ave.
   Monterey, CA 93943

4. Chairman, Computer Science — 1
   Naval Postgraduate School
   Code CS,
   833 Dyer Rd.
   Monterey, CA 93943-5118

5. Debra Hensgen — 1
   OpenTV, Inc.
   401 East Middlefield Road
   Mountain View, CA 94043

6. Rudy Darken — 1
   Naval Postgraduate School
   Code CS/DR, Computer Sciences Dept.
   833 Dyer Rd.
   Monterey, CA 93943-5118

7. Cynthia Irvine — 1
   Naval Postgraduate School
   Code CS/IC, Computer Sciences Dept.
   833 Dyer Rd.
   Monterey, CA 93943-5118

8. Lantz Moore — 1
   13249 Saddle Ridge Road
   Lakeside, CA 92040

9. Susie Tsai 1
5534, South, Oakhurst Place
Seattle, WA 98118

10. Pedro T.H. Tsai 1
979 Syida Drive
Pacific Grove, CA 93950

69 290NP6 TH 2935
6/02 22527-200 NLE

DUDLEY KNOX LIBRARY
3 2768 00410849 8